

Metrics for sparse graphs

Béla Bollobás^{*†} Oliver Riordan[‡]

December 14, 2008

Abstract

Recently, Bollobás, Janson and Riordan introduced a very general family of random graph models, producing inhomogeneous random graphs with $\Theta(n)$ edges. Roughly speaking, there is one model for each *kernel*, i.e., each symmetric measurable function from $[0, 1]^2$ to the non-negative reals, although the details are much more complicated, to ensure the exact inclusion of many of the recent models for large-scale real-world networks.

A different connection between kernels and random graphs arises in the recent work of Borgs, Chayes, Lovász, Sós, Szegedy and Vesztegombi. They introduced several natural metrics on dense graphs (graphs with n vertices and $\Theta(n^2)$ edges), showed that these metrics are equivalent, and gave a description of the completion of the space of all graphs with respect to any of these metrics in terms of *graphons*, which are essentially bounded kernels. One of the most appealing aspects of this work is the message that sequences of inhomogeneous quasi-random graphs are in a sense completely general: any sequence of dense graphs contains such a subsequence. Alternatively, their results show that certain natural models of dense inhomogeneous random graphs (one for each graphon) cover the space of dense graphs: there is one model for each point of the completion, producing graphs that converge to this point.

Our aim here is to briefly survey these results, and then to investigate to what extent they can be generalized to graphs with $o(n^2)$ edges. Although many of the definitions extend in a simple way, the connections between the various metrics, and between the metrics and random graph models, turn out to be much more complicated than in the dense case. We shall prove many partial results, and state even more conjectures and open problems, whose resolution would greatly enhance the currently rather unsatisfactory theory of metrics on sparse graphs. This paper deals mainly with graphs with $o(n^2)$ but $\omega(n)$ edges: a companion paper will discuss the (more problematic still) case of *extremely sparse* graphs, with $O(n)$ edges.

^{*}Department of Pure Mathematics and Mathematical Statistics, Wilberforce Road, Cambridge CB3 0WB, UK and Department of Mathematical Sciences, University of Memphis, Memphis TN 38152, USA. E-mail: b.bollobas@dpmps.cam.ac.uk.

[†]Research supported in part by NSF grants CCR-0225610 and DMS-0505550 and ARO grant W911NF-06-1-0076

[‡]Mathematical Institute, University of Oxford, 24–29 St Giles', Oxford OX1 3LB, UK. E-mail: riordan@maths.ox.ac.uk.

Contents

1	Introduction	2
2	Dense graphs	5
2.1	The subgraph distance	5
2.2	The cut distance	8
2.3	Kernels and (quasi-)random graphs	11
2.4	Equivalent kernels	12
3	Subgraph counts for sparse graphs	18
3.1	Bounded and unbounded kernels	21
3.2	Non-uniform random graphs	23
3.3	Subgraph counts in the uniform case	24
3.4	Partial results in the almost dense, uniform case	30
3.5	Extensions to lower densities.	34
4	Szemerédi’s Lemma and the cut metric	36
4.1	Weakly regular partitions	38
4.2	Strongly regular partitions	41
4.3	Szemerédi’s lemma and convergence in the cut norm	43
5	Comparison between cut and count convergence	47
5.1	Admissible subgraphs and their counts	48
5.2	Conjectured equivalence between cut and count convergence . . .	51
5.3	Partial results: embedding lemmas	55
5.4	Embeddings or homomorphisms?	66
6	The partition metric	73
6.1	Partition matrices and the partition metric	73
6.2	The relationship between the cut and partition metrics	76
7	Discussion and closing remarks	78
7.1	Models and metrics	78
7.2	Closing Remarks	80

1 Introduction

In recent years, much work has been done constructing and analyzing mathematical models of real-world networks. The random graphs in these models are inhomogeneous – in fact, many of them have degree sequences with power law distributions. In [8], Bollobás, Janson and Riordan defined a very general model of an n -vertex random graph $G(n, \kappa)$ with conditional independence between the edges which includes as special cases many of the models of real-world networks that have been studied, and proved numerous results about the random graphs generated by this model, including results about their component structure and

the point and nature of the phase transition in them. Here the *kernel* κ is a symmetric measurable function from $[0, 1]^2$ to $[0, \infty)$ satisfying some mild conditions. (Some of these conditions arise due to the very general nature of other parts of the model, and can be weakened in other contexts; see [9] and [10] for a discussion of this.) Just like the real-world graphs that motivated the construction of the BJR model, the random graphs $G(n, \kappa)$ are sparse in the sense that the expected number of edges is $O(n)$ (in fact, $(c + o(1))n$ for some constant c). In [8] the kernel κ was used to define a multi-type branching process \mathcal{X}_κ whose survival probability is closely related to the component structure of $G(n, \kappa)$.

In order to decide how well our random graph $G(n, \kappa)$ approximates a given real-world graph G_n , it would be desirable to establish a *distance* between a random graph model and a graph, so that the approximation is judged to be better and better as the distance tends to 0. Putting it slightly differently, we should like to define a metric on the set of sparse finite graphs so that a Cauchy sequence consists of graphs that are in some sense ‘similar’, and the limit of such a (not eventually constant) sequence is naturally identified with a suitable random graph model. For *dense* graphs, graphs with n vertices and at least cn^2 edges, such a program has been carried out very successfully in a series of papers by (various subsets of) Borgs, Chayes, Lovász, Sós, Szegedy and Vesztergombi (see [13, 14, 33, 34, 15, 16] and the references therein). In particular, they introduced several metrics on the space of dense finite (weighted) graphs and showed them to be equivalent. The limiting objects, i.e., the additional points in the completion, turn out to be *graphons*, that is, bounded symmetric measurable functions from $[0, 1]^2$ to \mathbb{R} . The corresponding random graph models, called *W-random graphs* in [34], are the natural dense version of $G(n, \kappa)$; see Subsection 2.3.

The only difference between kernels and graphons is that the latter are bounded, while the former must be allowed to be unbounded in order to model, for example, highly inhomogeneous real-world networks. In many fundamental questions (for example those concerning the phase transition), this difference is substantial. The appearance of graphons or kernels in the two different contexts described above suggests the existence of interesting connections between these areas. One such connection is described by Bollobás, Borgs, Chayes and Riordan [7], who study (sparse) random subgraphs of arbitrary dense graphs; this has recently been extended by Bollobás, Janson and Riordan [10].

We have several aims in this paper. First, we shall review some of the results of Borgs, Chayes, Lovász, Sós, Szegedy and Vesztergombi mentioned above. Our main aim is then to take the first tentative steps towards a general theory of metrics on sparse graphs; in particular, we shall investigate to what extent these ideas can be carried over to the sparse setting, and what can be said about the connection between the metrics and the ideas of Bollobás, Janson and Riordan. As we shall see, the difficulties that arise are considerably greater than in the dense case; in fact, the difficulties increase as the graphs get sparser. The *almost dense* case $e(G_n) = n^{2-o(1)}$ is already rather different from the dense case; the *extremely sparse* case $e(G_n) = \Theta(n)$, which will be studied in a companion paper [11], is *very* different indeed, having many novel features. We shall prove

numerous results, but the picture we obtain is much less complete than that obtained by Borgs et al in the dense case. In fact, perhaps our most important aim is to identify some of the main problems and conjectures whose resolution would enhance the theory of metrics on sparse graphs.

An important tool in the study of metrics on spaces of dense graphs is Szemerédi’s Regularity Lemma. While there is a version of Szemerédi’s Lemma for sparse graphs (with $o(n^2)$ but $\omega(n)$ edges) satisfying a mild additional condition, there is no satisfactory counting/embedding lemma for counting (or even finding) small subgraphs using regular partitions. This is one of the reasons why sparse graphs are much more difficult to handle than dense ones. One of our main aims is to prove such a counting lemma for certain subgraphs, greatly extending a result of Chung and Graham [17].

The rest of the paper is organized as follows. The next section is about dense graphs and kernels; we start by briefly recalling some of the definitions and results of Borgs, Chayes, Lovász, Sós, Szegedy and Vesztergombi whose generalization we shall discuss, focussing in particular on the cut metric. Then, in Subsection 2.4, we show that these results are closely connected to the question of when two kernels are ‘equivalent’; we shall need this notion of equivalence when we come to sparse graphs.

The rest of the paper concerns sparse graphs, i.e., graphs with n vertices and $o(n^2)$ edges: in Section 3 we consider subgraph counts in sparse (but mostly not too sparse) graphs, stating a conjecture that generalizes the main result of Lovász and Szegedy [34], and proving various partial results, concentrating especially on the *uniform* case, i.e., on sparse quasi-random graphs. In Section 4 we turn to Szemerédi’s Lemma for sparse graphs satisfying an appropriate ‘bounded density’ assumption, and the consequences for questions of convergence in the cut metric.

Sections 5 is the longest and most important section of the paper. In it we discuss the relationship between the cut metric and the count metric (to be defined) in the sparse case. As well as proposing various conjectures extending the results of Borgs, Chayes, Lovász, Sós and Vesztergombi, we prove several partial results, amounting to ‘sparse counting lemmas’ with various assumptions; these results, Theorem 5.14 and its variants Theorems 5.15 and 5.17, are the most substantial results in the paper.

In Section 6 we briefly discuss another metric considered by Borgs, Chayes, Lovász, Sós and Vesztergombi, showing that for graphs that are sparse, but not too sparse, it is equivalent to the cut metric. In the *extremely sparse* case, considering graphs with bounded average degree, the partition metric turns out to be much more useful than the cut metric. This and a discussion of the many problems and interesting open questions concerning metrics on extremely sparse graphs will be the topic of a companion paper [11].

In Section 7 we return briefly to the relationship between metrics and random graph models, and close with some final remarks summarizing our main results and conjectures.

Throughout the paper we use standard graph theoretic notation as in [4]. For example, $|G|$ and $e(G)$ denote respectively the number of vertices and number

of edges of a graph G .

2 Dense graphs

There are many natural definitions of what it means for two graphs to be ‘close’, and corresponding metrics and notions of Cauchy/fundamental sequences. These tend to be particularly natural for ‘dense’ graphs, with $\Theta(n^2)$ edges. Several of these metrics have been studied by Borgs, Chayes, Lovász, Sós and Vesztergombi [15, 16], who showed that they are equivalent, and that there is a natural completion of the space of graphs under any of these metrics. In this section we briefly recall some of these definitions and results; we are not aiming to give a comprehensive survey of the results of these papers, discussing only those that will be relevant for us here. Although most of the results mentioned in Subsections 2.1–2.3 will be from Lovász and Szegedy [34] and [15, 16], we shall not always adopt their notation or terminology, or indeed follow their definitions exactly.

Borgs, Chayes, Lovász, Sós and Vesztergombi [15, 16] consider *weighted graphs*, with weights on the edges and on the vertices. For the results we shall describe, this makes essentially no difference. In what follows, we consider only unweighted graphs; while much of what we shall say presumably carries over to suitably weighted graphs, the definitions for weighted graphs are not as natural in the sparse case, and are likely to introduce more additional complications than new insights.

2.1 The subgraph distance

The basic starting point is to consider, for each fixed graph F , the number of copies of F in a large graph G , i.e., the number $X_F(G)$ of subgraphs of G isomorphic to F . Recall that a *homomorphism* from a graph F to a graph G is a function $\phi : V(F) \rightarrow V(G)$ such that $\phi(x)\phi(y) \in E(G)$ whenever $xy \in E(F)$. Although $X_F(G)$ (for example, the number of triangles in G) is the most natural basic notion in this context, it turns out to be cleaner to work with $\text{emb}(F, G)$, the number of injective homomorphisms or *embeddings* of F into G . Note that

$$\text{emb}(F, G) = \text{aut}(F)X_F(G),$$

so $X_F(G)$ and $\text{emb}(F, G)$ contain the same information. Working with the latter avoids constant factors $\text{aut}(F)$ in many formulae.

If F has k vertices, then for $n \geq k$ we have $\text{emb}(F, K_n) = n_{(k)} = n(n-1) \cdots (n-k+1)$, so the natural normalization is to work with

$$s(F, G) = \frac{\text{emb}(F, G)}{n_{(k)}} = \frac{X_F(G)}{X_F(K_n)} \in [0, 1],$$

where, as usual, $n = |G|$ is the number of vertices of G . If $|F| > |G|$ then the above ratio is not defined, and we set $s(F, G) = 0$.

Let \mathcal{F} denote the set of isomorphism classes of finite graphs; sometimes it will be convenient to enumerate \mathcal{F} in an arbitrary way, writing $\mathcal{F} = \{F_1, F_2, \dots\}$. (More formally, we shall take each F_i to be a representative of an isomorphism class.) The graph parameters $s(F, \cdot)$, $F \in \mathcal{F}$, define a natural family of equivalent metrics on \mathcal{F} , by mapping \mathcal{F} into $[0, 1]^\infty$ (or into $[0, 1]^{\mathcal{F}}$). Indeed, for any finite graph G , set

$$s(G) = (s_i(G))_{i=1}^\infty \in [0, 1]^\infty,$$

where $s_i(G) = s(F_i, G)$. Let d be any metric on $X = [0, 1]^\infty$ which gives the product topology, for example $d(s, t) = \sum_{i=1}^\infty 2^{-i} |s_i - t_i|$. We may define the *subgraph distance* of two graphs G_1, G_2 as

$$d_{\text{sub}}(G_1, G_2) = d(s(G_1), s(G_2)).$$

It is easy to see that this defines a metric on \mathcal{F} : indeed, given $G \in \mathcal{F}$, among graphs F with $s(F, G) > 0$, there is a unique graph with $|F| + e(F)$ maximal, namely G . Thus the map $G \mapsto s(G)$ is injective. Furthermore, considering $s(E_{n+1}, G)$, where E_{n+1} is the empty graph with $n+1$ vertices, we see that the distance between any graph G with n vertices and the set of graphs with more than n vertices is positive. It follows that the metric space $(\mathcal{F}, d_{\text{sub}})$ is discrete.

A sequence (G_n) of graphs is Cauchy with respect to d_{sub} if and only if, for each $F \in \mathcal{F}$, the sequence $s(F, G)$ converges. Such sequences are sometimes called ‘convergent’, although they do not converge in the metric space $(\mathcal{F}, d_{\text{sub}})$. Note that if (G_n) is Cauchy then, since $(\mathcal{F}, d_{\text{sub}})$ is discrete, either (G_n) is eventually constant, or $|G_n| \rightarrow \infty$.

Many minor variations on the definition of d_{sub} are possible. For example, instead of considering the number of embeddings of F into G , one can consider the number $\text{hom}(F, G)$ of homomorphisms from F to G . If $|F| = k$ and $|G| = n$, then the number of non-injective homomorphisms from F to G is at most $\binom{k}{2} n^{k-1} = O(n^{k-1})$, so setting

$$t(F, G) = \text{hom}(F, G)/n^k$$

we have

$$t(F, G) = s(F, G) + O(n^{-1}) \tag{1}$$

for each F . Hence, in this dense case, the parameters $s(F, \cdot)$ and $t(F, \cdot)$ are essentially equivalent. [There is a minor difference that, working with homomorphisms, one ends up with a pseudo-metric: if G is any graph and $G^{(r)}$ is the *blow-up* of G obtained by making r copies of each vertex, joined to all copies of its neighbours, then $t(F, G^{(r)}) = t(F, G)$ for all $F \in \mathcal{F}$ and $r \geq 1$.] Also, one can pass easily back and forth between subgraph counts and counts of induced subgraphs using inclusion–exclusion.

One of the key properties of the metric d_{sub} is that there is a natural description of the (clearly compact) completion of $(\mathcal{F}, d_{\text{sub}})$, in terms of *standard kernels* (also called *graphons*). Here a *kernel* is a symmetric measurable function from $[0, 1]^2$ to $[0, \infty)$; a standard kernel is one taking values in $[0, 1]$. In other contexts, one considers more general bounded kernels, taking values in $[0, M]$ or

$[-M, M]$, $M > 0$, or general *signed kernels* taking values in \mathbb{R} . One can extend the definition of $s(F, G)$ (or of $t(F, G)$) to kernels in a natural way: given a finite graph F with vertex set $\{1, 2, \dots, k\}$, let

$$s(F, \kappa) = \int_{[0,1]^k} \prod_{ij \in E(F)} \kappa(x_i, x_j) \prod_{i=1}^k dx_i. \quad (2)$$

(Some authors use the notation $t(F, \kappa)$ for the same quantity.) This formula has a natural interpretation as the normalized ‘number’ of embeddings of F into a weighted graph with the uncountable vertex set $[0, 1]$, with edge weights given by κ . Of course, in this context there is no difference between embeddings and homomorphisms.

Lovász and Szegedy [34] proved (essentially) the following result.

Theorem 2.1. *Let (G_n) be a Cauchy sequence in $(\mathcal{F}, d_{\text{sub}})$. Then either (G_n) is eventually constant, or there is a standard kernel κ such that $s(F, G_n) \rightarrow s(F, \kappa)$.* \square

Let us remark that the result proved in [34] concerns t rather than s , which makes no difference, except that a separate case for eventually constant sequences is then not needed. Here, the distinction is informative: considering the parameters $s(E_k, G_n)$ for each k shows that in the second case above we have $|G_n| \rightarrow \infty$.

Of course, (2) allows one to extend the metric d_{sub} to standard kernels, obtaining in the first instance a pseudo-metric on the set of standard kernels. There is a natural notion of equivalence for kernels, which one can think of as a two dimensional version of the equivalence relation on random variables given by $X \sim Y$ if X and Y have the same distribution; the details are somewhat technical, and not essential for understanding the metrics discussed here, so we postpone them to Subsection 2.4. We write \sim for this relation, and \mathcal{K} for the set of equivalence classes of standard kernels under \sim . Borgs, Chayes and Lovász [12] have shown that $\kappa_1 \sim \kappa_2$ if and only if $d_{\text{sub}}(\kappa_1, \kappa_2) = 0$ (see also Theorem 2.8), so d_{sub} induces a metric on \mathcal{K} . The metric space $(\mathcal{K}, d_{\text{sub}})$ is complete (the result about Cauchy sequences of graphs above applies just as well to standard kernels). Hence, the completion of $(\mathcal{F}, d_{\text{sub}})$ is obtained by adding to \mathcal{F} the set \mathcal{K} of all equivalence classes of standard kernels, and using the map $s : \mathcal{F} \cup \mathcal{K} \rightarrow [0, 1]^\infty$ to extend d_{sub} to $\mathcal{F} \cup \mathcal{K}$.

There is a natural way to associate a standard kernel κ_G to a graph G with n vertices: divide $[0, 1]$ into n intervals I_1, \dots, I_n of equal length (we may and shall ignore the question of which endpoints are included), and set κ_G to be 1 on $I_i \times I_j$ if $ij \in E(G)$, and 0 otherwise. One slight advantage of using t rather than s is that

$$t(F, G) = s(F, \kappa_G)$$

for all graphs F and G . However, the metric obtained using t is only a pseudo-metric, since graphs on different numbers of vertices may correspond to the same kernel, for example if one is a blow-up of the other.

We say that a kernel κ is of *finite type* if there is a partition of $[0, 1]$ into measurable sets A_1, \dots, A_k so that κ is constant on each of the rectangles $A_i \times A_j$. Note that κ_G is always of finite type.

2.2 The cut distance

Borgs, Chayes, Lovász, Sós and Vesztergombi [15] considered another natural metric on graphs or kernels, namely, the *cut metric*, based on a norm used by Frieze and Kannan [23]. For any integrable function $\kappa : [0, 1]^2 \rightarrow \mathbb{R}$, its *cut norm* $\|\kappa\|_{\text{cut}}$ is defined by

$$\|\kappa\|_{\text{cut}} = \sup_{S, T \subset [0, 1]} \left| \int_{S \times T} \kappa(x, y) dx dy \right|, \quad (3)$$

where the supremum is over all pairs of measurable subsets of $[0, 1]$. It is easily seen that this defines a norm on $L^\infty([0, 1]^2)$. In fact, there are several variations of this definition: one can take

$$\|\kappa\|_{\text{cut}} = \sup_{S \subset [0, 1]} \left| \int_{S \times S^c} \kappa(x, y) dx dy \right|, \quad (4)$$

where $S^c = [0, 1] \setminus S$, or one can take the supremum in (3) only over sets S, T with $S \cap T = \emptyset$. It is easy to check that these variations only affect the norm up to an (irrelevant) constant factor (see [15]), so we shall feel free to use whichever definition is most convenient in any given context.

There is yet another definition of $\|\kappa\|_{\text{cut}}$ that is more natural from the point of view of functional analysis, namely

$$\|\kappa\|_{\text{cut}} = \sup_{\|f\|_\infty, \|g\|_\infty \leq 1} \int_{[0, 1]^2} \kappa(x, y) f(x) g(y) dx dy,$$

where the supremum is taken over all pairs of measurable functions from $[0, 1]$ to $[-1, +1]$. Since the integral above is linear with respect to each of f and g , the supremum is attained at some functions taking values in $\{-1, +1\}$, and it follows immediately that this version of the cut norm is again within a constant factor of that defined by (3). As noted in [10], for example, this last definition is the most natural from the point of view of functional analysis: it is the dual of the projective tensor product norm in $L^\infty \hat{\otimes} L^\infty$, and is thus the injective tensor product norm in $L^1 \hat{\otimes} L^1$. Equivalently, this is just the norm of the integral operator with kernel κ , treated as a map from L^∞ to L^1 .

Before turning to the cut metric we need one further definition. Given a kernel κ and a measure-preserving map $\tau : [0, 1] \rightarrow [0, 1]$, let $\kappa^{(\tau)}$ be the kernel defined by

$$\kappa^{(\tau)}(x, y) = \kappa(\tau(x), \tau(y)). \quad (5)$$

If τ is a bijection, then we call τ a *rearrangement* of $[0, 1]$, and $\kappa^{(\tau)}$ a *rearrangement* of κ . (It is perhaps more natural to consider measure-preserving bijections between two subsets of $[0, 1]$ with measure 1; this makes no difference.) Two

kernels κ_1 and κ_2 are *naively equivalent* if one is a rearrangement of the other, more precisely, if there is a rearrangement τ of $[0, 1]$ such that

$$\kappa_1(x, y) = \kappa_2^{(\tau)}(x, y) \quad \text{for a.e. } (x, y) \in [0, 1]^2. \quad (6)$$

In this case we write $\kappa_1 \approx \kappa_2$, noting that \approx is an equivalence relation.

The *cut metric* d_{cut} on the set of standard kernels may be defined as follows:

$$d_{\text{cut}}(\kappa_1, \kappa_2) = \inf_{\kappa'_2 \approx \kappa_2} \|\kappa_1 - \kappa'_2\|_{\text{cut}}. \quad (7)$$

Clearly, this defines a pseudo-metric on standard kernels; in particular, if $\kappa_1 \approx \kappa_2$, then $d_{\text{cut}}(\kappa_1, \kappa_2) = 0$. The reverse implication does not hold; in fact, $d_{\text{cut}}(\kappa_1, \kappa_2) = 0$ if and only if $\kappa_1 \sim \kappa_2$, where \sim is the equivalence relation to be defined in Subsection 2.4. Hence, d_{cut} induces a metric on the set \mathcal{K} of equivalence classes of standard kernels under the relation \sim .

As noted above, there is a standard kernel κ_G naturally associated to each graph G , although the map $G \mapsto \kappa_G$ from \mathcal{F} to \mathcal{K} is not injective. One extends the cut metric to a pseudo-metric on graphs by setting

$$d_{\text{cut}}(G_1, G_2) = d_{\text{cut}}(\kappa_{G_1}, \kappa_{G_2}), \quad (8)$$

and to $\mathcal{F} \cup \mathcal{K}$ similarly.

For graphs G_1, G_2 on n vertices, there is a much more natural variant of their cut distance: let $\widehat{d}_{\text{cut}}(G_1, G_2)$ be the smallest ε for which we can identify the vertices of G_1 with those of G_2 such that for any bipartition of the vertex set, the corresponding cuts in G_1 and G_2 have sizes within εn^2 . In terms of kernels,

$$\widehat{d}_{\text{cut}}(G_1, G_2) = \min_{\kappa \approx_n \kappa_{G_2}} \|\kappa_{G_1} - \kappa\|_{\text{cut}}, \quad (9)$$

where $\kappa_1 \approx_n \kappa_2$ if (6) holds for some map τ that simply permutes the intervals I_n corresponding to the vertices, and we take (4) as the definition of the cut norm. Note that the supremum implied by (4) in the definition (9) is over all bipartitions of $[0, 1]$, not just those corresponding to bipartitions of the vertices; it is very easy to see that this makes no difference: the supremum is attained at a vertex bipartition.

Comparing (8) and (9), since the infimum in the former is taken over a larger set, one trivially has $d_{\text{cut}}(G_1, G_2) \leq \widehat{d}_{\text{cut}}(G_1, G_2)$. Borgs, Chayes, Lovász, Sós and Vesztergombi [15] noted that strict inequality is possible. For example, taking (4) as the definition of the cut norm, let G_1 be a triangle, and let G_2 be the graph with 3 vertices and one edge. For any pairing of the vertices of G_1 with those of G_2 , the ‘worst’ cut is the one in which the isolated vertex of G_2 is placed into one part and the other two vertices into the other part. This cut has 2 edges in G_1 but no edges in G_2 , so $\widehat{d}_{\text{cut}}(G_1, G_2) = 2/9$. On the other hand, consider the blow-ups $G_1^{(2)}$, a complete tripartite graph with two vertices in each class, and $G_2^{(2)}$, a C_4 with two isolated vertices added. Pairing the vertices of $G_1^{(2)}$ and $G_2^{(2)}$ by placing two opposite vertices of the C_4 in one

class of $G_1^{(2)}$, and the other vertices in different classes, we realize $G_2^{(2)}$ as a subgraph of $G_1^{(2)}$ in such a way that the 8 edges of $G_1^{(2)}$ not present in $G_2^{(2)}$ form a non-bipartite graph, so *every* cut cuts at most 7 of these extra edges. It follows that $\widehat{d}_{\text{cut}}(G_1^{(2)}, G_2^{(2)}) \leq 7/6^2$. In fact, one can check that with the vertices paired in this way the maximum difference between the sizes of corresponding cuts in $G_1^{(2)}$ and $G_2^{(2)}$ is 6, so

$$d_{\text{cut}}(G_1, G_2) \leq \widehat{d}_{\text{cut}}(G_1^{(2)}, G_2^{(2)}) \leq 6/6^2 = 1/6 < 2/9 = \widehat{d}_{\text{cut}}(G_1, G_2),$$

showing that d_{cut} and \widehat{d}_{cut} do not always agree. For questions of convergence, however, the two metrics are equivalent: as shown in [15],

$$d_{\text{cut}}(G_1, G_2) \leq \widehat{d}_{\text{cut}}(G_1, G_2) \leq 32d_{\text{cut}}(G_1, G_2)^{1/67}.$$

At first sight it is not clear why the cut metric should be interesting: after all, what is the significance of two graphs having almost the same number of edges in all corresponding cuts? One very important consequence of this property is that their subgraph counts are close, as shown by the following simple lemma from Borgs, Chayes, Lovász, Sós and Vesztegombi [15].

Lemma 2.2. *Let κ and κ' be two standard kernels. Then for every graph F we have*

$$|s(F, \kappa) - s(F, \kappa')| \leq e(F) \|\kappa - \kappa'\|_{\text{cut}}.$$

Proof. Before we embark on the proof, we extend the definition of $s(F, \kappa)$ slightly. Fix the graph F , taking its vertex set to be $[k] = \{1, 2, \dots, k\}$, as usual, and list the edges of F as $\{i_1 j_1, \dots, i_m j_m\}$. Given a sequence $(\kappa_1, \dots, \kappa_m)$ of standard kernels, set

$$s(F; \kappa_1, \dots, \kappa_m) = \int_{[0,1]^k} \prod_{r=1}^m \kappa_r(x_{i_r}, x_{j_r}) \prod_{i=1}^k dx_i.$$

Thus $s(F, \kappa) = s(F; \kappa, \dots, \kappa)$. We claim that for any graph F with m edges and any standard kernels $\kappa_1, \kappa_2, \dots, \kappa_m$ and κ'_1 , we have

$$|s(F; \kappa_1, \kappa_2, \dots, \kappa_m) - s(F; \kappa'_1, \kappa_2, \dots, \kappa_m)| \leq \|\kappa_1 - \kappa'_1\|_{\text{cut}}. \quad (10)$$

Applying this $m = e(F)$ times, changing one kernel from κ to κ' each time, the lemma follows.

It remains to prove (10), which is easy. Suppose without loss of generality that the first edge is 12, so $i_1 = 1$ and $j_1 = 2$. Our task is to bound

$$\Delta = \int_{[0,1]^k} (\kappa_1(x_1, x_2) - \kappa'_1(x_1, x_2)) \prod_{r=2}^m \kappa_r(x_{i_r}, x_{j_r}) \prod_{i=1}^k dx_i$$

Collecting the terms in the product that involve x_1 or x_2 , we may write this product as $f_0(\mathbf{x})f_1(x_1, \mathbf{x})f_2(x_2, \mathbf{x})$, where $\mathbf{x} = (x_3, \dots, x_k)$ and each f_i (being a product of standard kernels evaluated at certain places) takes values in

$[0, 1]$. Now from (3), it is immediate that if f and g take values in $[0, 1]$, then $|\int \kappa(x, y) f(x) g(y) dx dy| \leq \|\kappa\|_{\text{cut}}$. Applying this with \mathbf{x} fixed, and then integrating over \mathbf{x} , it follows that $|\Delta| \leq \|\kappa_1 - \kappa'_1\|_{\text{cut}}$, as required. \square

Corollary 2.3. *Let (G_n) be a sequence of graphs with $|G_n| \rightarrow \infty$, and let κ be a standard kernel. If $d_{\text{cut}}(G_n, \kappa) \rightarrow 0$ then $d_{\text{sub}}(G_n, \kappa) \rightarrow 0$.*

Proof. Let $\kappa_n = \kappa_{G_n}$, so by definition $d_{\text{cut}}(G_n, \kappa) = d_{\text{cut}}(\kappa_n, \kappa)$. By Lemma 2.2, for every F we have $s(F, \kappa_n) \rightarrow s(F, \kappa)$. But $s(F, \kappa_n) = t(F, G_n)$, while from (1) we have $s(F, G_n) = t(F, G_n) + o(1)$. Thus $s(F, G_n) \rightarrow s(F, \kappa)$ for each F , i.e., $d_{\text{sub}}(G_n, \kappa) \rightarrow 0$. \square

We have just seen that convergence in d_{cut} implies convergence in d_{sub} ; one of the main results of Borgs, Chayes, Lovász, Sós and Vesztergombi, namely Theorem 2.6 in [15], gives a converse of this. This result states that the metrics d_{sub} (defined using t rather than s) and d_{cut} are equivalent, in the sense that (G_n) is a Cauchy sequence for d_{sub} if and only if it is a Cauchy sequence for d_{cut} . In the light of the various other results of Lovász and Szegedy [34] and Borgs, Chayes, Lovász, Sós and Vesztergombi [15], this statement may be reformulated in our notation as follows.

Theorem 2.4. *Let (G_n) be a sequence of graphs or standard kernels with $|G_n| \rightarrow \infty$, where we take $|G_n| = \infty$ if G_n is a kernel, and let κ be a standard kernel. Then $d_{\text{sub}}(G_n, \kappa) \rightarrow 0$ if and only if $d_{\text{cut}}(G_n, \kappa) \rightarrow 0$.* \square

An immediate consequence of this result is the following, Corollary 3.10 in [15].

Corollary 2.5. *Let κ and κ' be two bounded kernels. Then $s(F, \kappa) = s(F, \kappa')$ for every F if and only if $d_{\text{cut}}(\kappa, \kappa') = 0$.* \square

We shall return to a discussion of kernels at cut distance 0 shortly.

2.3 Kernels and (quasi-)random graphs

As well as going from graphs to kernels, one can go from kernels to *random* graphs in a very natural way, as in Section 2.6 of Lovász and Szegedy [34], or as in Bollobás, Janson and Riordan [8] for the sparse case. Indeed, given a standard kernel κ and an $n \geq 1$, let $G(n, \kappa)$ be the random graph on $[n]$ defined as follows: first let x_1, \dots, x_n be iid with the uniform distribution on $[0, 1]$. Given the x_i , join each pair of vertices independently, joining i and j with probability $\kappa(x_i, x_j)$. The resulting graph is called a κ -*random graph* by Lovász and Szegedy [34], although they use W as their default symbol for a kernel. It is easy to check, for example by the second moment method, that, for each F , the random variable $s(F, G(n, \kappa))$ converges (in probability and in fact almost surely) to $s(F, \kappa)$ as $n \rightarrow \infty$. Thus the sequence $G(n, \kappa)$ converges almost surely to κ in the metric d_{sub} or d_{cut} . Note that if κ is constant and takes the value p , then we recover the usual Erdős–Rényi model $G(n, p)$: no confusion should

arise between the notation for the two models. (In fact, it was Gilbert [25] who introduced $G(n, p)$, while Erdős and Rényi [21] introduced a model, $G(n, m)$, that is essentially equivalent for many purposes. Since it was they who founded the theory of random graphs, both models are often referred to as Erdős–Rényi models.)

It is natural to view a sequence (G_n) converging to κ in d_{sub} as a sequence of ‘inhomogeneous quasi-random graphs’: when κ is constant, the convergence condition is equivalent to the standard notion of quasi-randomness, introduced by Thomason [37] in 1987 (although he called it pseudo-randomness) and studied in great detail by Chung, Graham and Wilson [18] and many others. The convergence of $G(n, \kappa)$ to κ in d_{sub} establishes that sequences generated by the natural inhomogeneous random model are also quasi-random, as one would hope. One of the most pleasing features of this whole subject area is the interpretation that inhomogeneous quasi-random graphs are completely general: any sequence of (dense) graphs has such a subsequence.

To take an alternative viewpoint, we may think of standard kernels as uncountable infinite graphs, and a ‘typical’ random graph $G(n, \kappa)$ as a good finite approximation to κ . Then the completion of \mathcal{F} is obtained by adding these infinite graphs, and the approximations $G(n, \kappa)$ (n large) are examples of finite graphs close to a given infinite graph. Taking this viewpoint it is natural *not* to identify a finite graph with a kernel. For another, slightly different, point of view, see Diaconis and Janson [19], where connections to certain infinite random graphs are described.

2.4 Equivalent kernels

In the light of Corollary 2.5, it is clearly important to understand which pairs of kernels have $d_{\text{cut}}(\kappa_1, \kappa_2) = 0$; this is also important for understanding d_{cut} itself. Fortunately, it turns that there is a natural notion of equivalence for kernels which gives the answer. Since this topic is only touched on in passing in Borgs, Chayes, Lovász, Sós and Vesztergombi [15], we shall go into some detail here.

Roughly speaking, we would like to say that two kernels are equivalent if one is obtained from the other simply by relabelling the ‘types’ in $[0, 1]$. It would seem that the notion \approx of naive equivalence defined in (6) is thus the right one, but a little thought shows that this is not the case; for this, the random viewpoint is very helpful.

So far, as in [15], we defined kernels only on $[0, 1]^2$. In view of the connection to random graphs discussed in the previous subsection, it is *a priori* more natural to work with a general probability space $(\Omega, \mathcal{F}, \mu)$ rather than $[0, 1]$ with Lebesgue measure, defining a standard kernel as a symmetric measurable function from the square of a probability space to $[0, 1]$. (This is the approach taken in the sparse case by Bollobás, Janson and Riordan [8].) However, almost all the time, we shall consider only kernels on $[0, 1]$; there are two reasons for doing so: firstly, graphs with n vertices correspond to kernels on the discrete space with n equiprobable elements, and $[0, 1]$ is the natural limit of these spaces. Secondly,

all probability spaces that one would ever wish to work with (all so-called ‘standard’ probability spaces) are isomorphic to Lebesgue measure on an interval, combined with (possibly) a finite or countable number of atoms. When studying kernels, the presence of atoms makes no difference: for example, a kernel on a finite measure space corresponds in a natural way to a piecewise constant kernel on $[0, 1]$. Hence it makes very good sense to consider only kernels on $[0, 1]$. For a formal reduction to the case of kernels on $[0, 1]$ in the context of random graphs, see Janson [27].

We may think of kernels as two-dimensional versions of random variables (not to be confused with vector valued random variables). Two random variables are equivalent if they have the same distribution. Equivalently, they are equivalent if they may be coupled so as to agree with probability 1. This is the definition we shall use for kernels.

Working, for the moment, on general (standard) probability spaces, and suppressing the σ -field of measurable sets in the notation, let (Ω_1, μ_1) and (Ω_2, μ_2) be two probability spaces. A *coupling* of (Ω_1, μ_1) and (Ω_2, μ_2) is simply a probability space (Ω, μ) together with measure-preserving maps $\sigma_i : \Omega \rightarrow \Omega_i$, $i = 1, 2$. Thus, if X is a uniformly random point of (Ω, μ) , then $\sigma_1(X)$ and $\sigma_2(X)$ are uniform on (Ω_1, μ_1) and (Ω_2, μ_2) , respectively. Let κ_i be a kernel on (Ω_i, μ_i) , $i = 1, 2$. Then κ_1 and κ_2 are *equivalent* if there is a coupling (Ω, μ) of the underlying probability spaces such that

$$\kappa_1(\sigma_1(x), \sigma_1(y)) = \kappa_2(\sigma_2(x), \sigma_2(y)) \text{ for } (\mu \times \mu)\text{-almost every } (x, y) \in \Omega^2.$$

In other words, extending the notation in (5) to arbitrary spaces, we require $\kappa_1^{(\sigma_1)} = \kappa_2^{(\sigma_2)}$ a.e.; we write \sim for the corresponding relation. Although this definition may seem a little complicated, as explained above it is in fact very natural.

Note that $\kappa_1 \approx \kappa_2$ implies $\kappa_1 \sim \kappa_2$: if $\kappa_1 = \kappa_2^{(\tau)}$, then one couples $x \in [0, 1] = \Omega_1$ with $\tau(x) \in \Omega_2$. (More formally, we may take $\Omega = \Omega_1$, with σ_1 the identity and $\sigma_2 = \tau$.) It is easy to see that the reverse implication does not hold: for example, consider the random variables Λ_1, Λ_2 on $[0, 1]$ given by $\Lambda_1(x) = x$ and $\Lambda_2(x) = 2x - \lfloor 2x \rfloor$; these both have the uniform distribution, but since one is 1-to-1 and the other 2-to-1, there is no measure-preserving bijection from one ground space to the other transforming one into the other. Setting $\kappa_i(x, y) = \Lambda_i(x)\Lambda_i(y)$, one obtains kernels with $\kappa_1 \sim \kappa_2$ but $\kappa_1 \not\approx \kappa_2$. (Recently, Borgs, Chayes and Lovász [12] have shown that if one excludes this phenomenon of ‘twins’, then \sim and \approx are equivalent; we refer the reader there for a precise statement.)

Returning to the special case of kernels on $[0, 1]$, essentially equivalent to the general case, couplings have a very simple description. All that matters is that, for a uniform point X of (Ω, μ) , the distribution of $(\sigma_1(X), \sigma_2(X))$ should have uniform marginals. Thus, couplings correspond to *doubly stochastic measures*, i.e., Borel measures μ on $[0, 1]^2$ with both marginals Lebesgue measure. In other words, we have $\kappa_1 \sim \kappa_2$ if and only if there is a doubly stochastic measure μ

such that

$$\kappa_1(x, y) = \kappa_2(u, v) \text{ for } (\mu \times \mu)\text{-a.e. } (x, u, y, v) \in [0, 1]^4. \quad (11)$$

At first sight, $[0, 1]^2$ is the most natural space to use to couple two kernels on $[0, 1]$, but there is another natural choice. Since $[0, 1]^2$ is isomorphic as a probability space to $[0, 1]$, we may construct the coupling on $[0, 1]$! Hence, $\kappa_1 \sim \kappa_2$ if and only if there are measure preserving maps $\sigma_1, \sigma_2 : [0, 1] \rightarrow [0, 1]$ such that $\kappa_1^{(\sigma_1)} = \kappa_2^{(\sigma_2)}$ for (Lebesgue) a.e. $(x, y) \in [0, 1]^2$. Putting this a little more symmetrically, we see that $\kappa_1 \sim \kappa_2$ if and only if

$$\exists \kappa, \sigma_1, \sigma_2 \text{ such that } \kappa = \kappa_1^{(\sigma_1)} \text{ a.e. and } \kappa = \kappa_2^{(\sigma_2)} \text{ a.e.}, \quad (12)$$

where κ is a kernel on $[0, 1]$ and σ_1 and σ_2 are measure-preserving maps from $[0, 1]$ to itself. Note that $\kappa \sim \kappa^{(\sigma)}$ for any kernel κ on $[0, 1]$ and any measure-preserving map from $[0, 1]$ to itself.

Since couplings rather than rearrangements give the proper notion of equivalence for two kernels, it is natural to use couplings rather than rearrangements in the definition of the cut metric. Indeed, Borgs, Chayes, Lovász, Sós and Vesztegombi [15] define the cut metric on standard (or simply bounded) kernels as follows:

$$d_{\text{cut}}(\kappa_1, \kappa_2) = \inf_{\mu \in \mathcal{M}} \sup_{S, T} \left| \int_{S \times T} (\kappa_1(x, y) - \kappa_2(u, v)) d\mu(x, u) d\mu(y, v) \right|, \quad (13)$$

where \mathcal{M} is the set of doubly stochastic measures on $[0, 1]^2$, S and T run over measurable subsets of $[0, 1]^2$, and the integral is over $(x, u) \in S$ and $(y, v) \in T$. As shown in [15], the definitions (7) and (13) coincide. (This is not hard to see – either formula defines a function that is continuous, indeed Lipschitz with constant 1, with respect to the cut norm, and hence continuous with respect to the L^1 norm. Since the finite-type kernels are dense in L^1 , it suffices to check the equality of the two definitions for finite-type kernels, which is straightforward. For the details, see [15].) Since (7) is much easier to work with than (13), we shall take the former as our definition of d_{cut} .

Although (7) is more convenient, there is a sense in which (13) is the ‘right’ definition. For example, as we shall now show, the infimum in (13) is always attained, unlike that in (7). This is not discussed in [15], where it is of no particular significance. Here, as in the bulk of the paper, unless otherwise specified, all kernels are kernels on $[0, 1]$, i.e., symmetric Lebesgue-measurable functions from $[0, 1]^2 \rightarrow [0, \infty)$. As noted above, it always suffices to consider kernels on $[0, 1]$. Recall that we call a kernel *standard* if it takes values in $[0, 1]$.

Lemma 2.6. *Let κ_1 and κ_2 be two standard kernels. Then there is a doubly stochastic measure μ achieving the infimum in (13).*

Proof. For $\mu \in \mathcal{M}$ set

$$d_\mu(\kappa_1, \kappa_2) = \sup_{S, T} \left| \int_{S \times T} (\kappa_1(x, y) - \kappa_2(u, v)) d\mu(x, u) d\mu(y, v) \right|, \quad (14)$$

so our aim is to show that $\inf_{\mu \in \mathcal{M}} d_\mu(\kappa_1, \kappa_2)$ is attained. Before doing so, let us note that in the supremum one may restrict the sets S and T in (14) to ‘nice’ sets. Let \mathcal{D} denote the set of finite unions of products of (half-open) intervals. Since μ is a finite Borel measure, for any measurable $S, T \subset [0, 1]^2$ and any $\varepsilon > 0$, there are sets $S', T' \in \mathcal{D}$ with $\mu(S \Delta S'), \mu(T \Delta T') < \varepsilon$. Since $\kappa_1 - \kappa_2$ is bounded by ± 1 , replacing S, T by S' and T' changes the value of the integral by at most 2ε . It follows that the supremum in (14) may be taken over $S, T \in \mathcal{D}$ without changing its value, as claimed.

It is well known that \mathcal{M} is (sequentially) compact in the topology in which $\mu_n \rightarrow \mu$ if and only if $\mu_n(A) \rightarrow \mu(A)$ for every set $A \in \mathcal{D}$. Indeed, writing \mathcal{D}_0 for the set of products of intervals with rational endpoints, since \mathcal{D}_0 is countable any sequence in \mathcal{M} has a subsequence (μ_n) such that $(\mu_n(A))$ converges for all $A \in \mathcal{D}_0$. Using the doubly stochastic property to bound the measure of a rectangle with one or more short sides, convergence for all $A \in \mathcal{D}$ follows easily, and one can check that the limiting values do define a measure μ . Note that one cannot require $\mu_n(A) \rightarrow \mu(A)$ for every measurable A : it is easy to construct sequences where μ is concentrated on, for example, the diagonal $S = \{(x, x)\}$, with $\mu_n(S) = 0$ for every n .

Let (μ_n) be a sequence of doubly stochastic measures for which $d_{\mu_n}(\kappa_1, \kappa_2) \rightarrow d_{\text{cut}}(\kappa_1, \kappa_2)$; such a sequence exists by the definition (13) of $d_{\text{cut}}(\kappa_1, \kappa_2)$. From the remark above, (μ_n) has a subsequence converging to some $\mu \in \mathcal{M}$ in the appropriate topology. Restricting to this subsequence, we may assume that $\mu_n(A) \rightarrow \mu(A)$ for every $A \in \mathcal{D}$.

Let $S = S_1 \times S_2$ and $T = T_1 \times T_2$, where S_1, S_2, T_1 and T_2 are all intervals in $[0, 1]$. We claim that

$$\int_{S \times T} \kappa(x, y) d\mu_n(x, u) d\mu_n(y, v) \rightarrow \int_{S \times T} \kappa(x, y) d\mu(x, u) d\mu(y, v) \quad (15)$$

as $n \rightarrow \infty$, for any standard kernel κ . Before proving this, let us show that the lemma follows.

For any $\nu \in \mathcal{M}$, let

$$f(S, T, \nu) = \int_{S \times T} (\kappa_1(x, y) - \kappa_2(u, v)) d\nu(x, u) d\nu(y, v),$$

so $d_\nu(\kappa_1, \kappa_2) = \sup_{S, T} |f(S, T, \nu)|$. Applying (15) with $\kappa = \kappa_1$ and $\kappa = \kappa_2$, we see that $f(S, T, \mu_n) \rightarrow f(S, T, \mu)$ holds whenever S and T are products of intervals. By additivity, it thus holds whenever S and T are in \mathcal{D} . Since $d_{\mu_n}(\kappa_1, \kappa_2) = \sup_{S, T} |f(S, T, \mu_n)|$, for $S, T \in \mathcal{D}$ we thus have

$$f(S, T, \mu) = \liminf f(S, T, \mu_n) \leq \liminf d_{\mu_n}(\kappa_1, \kappa_2) = d_{\text{cut}}(\kappa_1, \kappa_2).$$

As noted earlier, when defining $d_\mu(\kappa_1, \kappa_2) = \sup_{S, T} |f(S, T, \mu)|$, we may take the supremum instead over $S, T \in \mathcal{D}$, so it follows that $d_\mu(\kappa_1, \kappa_2) \leq d_{\text{cut}}(\kappa_1, \kappa_2)$. Since $d_{\text{cut}}(\kappa_1, \kappa_2) = \inf_{\mu' \in \mathcal{M}} d_{\mu'}(\kappa_1, \kappa_2)$, this infimum is attained (at μ), as claimed.

It remains to prove (15). But this is easy: for any interval $I \subset [0, 1]$, let μ_n^I be the measure on $[0, 1]$ defined by

$$\mu_n^I(A) = \mu_n(A \times I),$$

and define μ^I from μ similarly. Recall that $\mu_n \rightarrow \mu$ on products of intervals. Thus $\mu_n^I(A) \rightarrow \mu^I(A)$ whenever A is an interval, and hence whenever A is a finite union of intervals. Since $\mu_n, \mu \in \mathcal{M}$, we have that $\mu_n(A)$ and $\mu(A)$ are both at most the Lebesgue measure of A . It follows that $\mu_n^I(A) \rightarrow \mu^I(A)$ for any measurable $A \subset [0, 1]$, since for any ε we can approximate A by a finite union of intervals A' whose symmetric difference from A has Lebesgue measure at most ε . It also follows that if I and J are two intervals, and $A \subset [0, 1]^2$ is Lebesgue measurable, then

$$(\mu_n^I \times \mu_n^J)(A) \rightarrow (\mu^I \times \mu^J)(A).$$

Indeed, this follows by approximating A by a finite union of products of intervals. Considering level sets, we see that

$$\int f(x, y) d\mu_n^I(x) d\mu_n^J(y) \rightarrow \int f(x, y) d\mu^I(x) d\mu^J(y)$$

for any bounded measurable function f . Taking $f = \kappa(x, y)1_{x \in S_1}1_{y \in T_1}$, $I = S_2$ and $J = T_2$, this is exactly (15), completing the proof. \square

The special case of Lemma 2.6 where the distance is 0 is of particular interest.

Corollary 2.7. *Let κ_1 and κ_2 be two standard kernels. Then $d_{\text{cut}}(\kappa_1, \kappa_2) = 0$ if and only if $\kappa_1 \sim \kappa_2$.*

Proof. Using (13) as the definition of d_{cut} , if $\kappa_1 \sim \kappa_2$ then we certainly have $d_{\text{cut}}(\kappa_1, \kappa_2) = 0$; see (11).

Suppose then that $d_{\text{cut}}(\kappa_1, \kappa_2) = 0$. From Lemma 2.6, there is a $\mu \in \mathcal{M}$ such that $d_\mu(\kappa_1, \kappa_2) = 0$. Let ν be the signed measure on $[0, 1]^4$ defined by

$$d\nu(x, u, y, v) = (\kappa_1(x, y) - \kappa_2(u, v)) d\mu(x, u) d\mu(y, v).$$

Then $d_\mu(\kappa_1, \kappa_2) = 0$ says exactly that $\nu(S \times T) = 0$ for all measurable $S, T \subset [0, 1]^2$. Since ν is a signed Borel measure, it follows immediately that ν is the zero measure. Equivalently, $\kappa_1(x, y) - \kappa_2(u, v) = 0$ for $(\mu \times \mu)$ -a.e. points (x, u, y, v) . Referring to (11) again, we see that $\kappa_1 \sim \kappa_2$. \square

As we have seen, Corollary 2.7 is a simple exercise in measure theory. Using this corollary, and the equivalence of d_{cut} and d_{sub} proved by Borgs, Chayes, Lovász, Sós and Vesztegombi [15], one obtains the following characterization of equivalent (standard) kernels.

Theorem 2.8. *Let κ_1 and κ_2 be two standard kernels. Then $s(F, \kappa_1) = s(F, \kappa_2)$ holds for every finite graph F if and only if $\kappa_1 \sim \kappa_2$.*

Proof. Immediate from Corollaries 2.5 and 2.7. \square

The analogue of Theorem 2.8 for general (i.e., unbounded) kernels is false, even for ‘rank 1’ kernels with all counts $s(F, \kappa)$ finite. Indeed, if $\kappa(x, y) = f(x)f(y)$ for some $f : [0, 1] \rightarrow [0, \infty)$, then the quantities $s(F, \kappa)$ are easily seen to be products of moments of f , viewed as a random variable. As is well known, there are non-negative random variables with the same finite moments but different distributions; using two such random variables, one can construct non-equivalent unbounded kernels κ_1, κ_2 with $s(F, \kappa_1) = s(F, \kappa_2) < \infty$ for all F .

We have shown that it is not hard to deduce Theorem 2.8 from Theorem 2.4. In fact, these results are equivalent! The reverse implication is actually much easier.

Theorem 2.8 \implies Theorem 2.4. We write out the argument for a sequence of graphs; the treatment for kernels is essentially the same. Let (G_n) be a sequence of graphs with $|G_n| \rightarrow \infty$, and let κ be a standard kernel. From Corollary 2.3, if $d_{\text{cut}}(G_n, \kappa) \rightarrow 0$, then $d_{\text{sub}}(G_n, \kappa) \rightarrow 0$; it remains to prove the reverse implication.

As shown by Lovász and Szegedy [34] (see their Lemmas 5.1 and 5.2), repeatedly applying even the weak Frieze–Kannan [23] form of Szemerédi’s Lemma, it is easy to prove that any sequence (G_n) with $|G_n| \rightarrow \infty$ has a subsequence converging in d_{cut} to some standard kernel κ' . We shall not give the details of this argument here as we shall prove a corresponding statement in a more general setting in Corollary 4.7.

Suppose then that $d_{\text{sub}}(G_n, \kappa) \rightarrow 0$. Then by the observation above there is a subsequence (G_{n_k}) that converges in d_{cut} to some standard kernel κ' . But then, by Corollary 2.3, we have $d_{\text{sub}}(G_{n_k}, \kappa') \rightarrow 0$. Since $d_{\text{sub}}(G_n, \kappa) \rightarrow 0$ we must have $d_{\text{sub}}(\kappa, \kappa') = 0$, i.e., $s(F, \kappa) = s(F, \kappa')$ for all F . Thus, by Theorem 2.8, we have $\kappa \sim \kappa'$, so $d_{\text{cut}}(\kappa, \kappa') = 0$. Thus $d_{\text{cut}}(G_{n_k}, \kappa) \rightarrow 0$.

We have shown that (G_n) has a subsequence converging to κ in d_{cut} . This argument applies equally well to any subsequence of (G_n) , and it follows immediately that the whole sequence converges, i.e., $d_{\text{cut}}(G_n, \kappa) \rightarrow 0$, as required. \square

As we have just seen, Theorem 2.4, one of the main results of Borgs, Chayes, Lovász, Sós and Vesztergombi [15], is equivalent to Theorem 2.8. As far as we are aware, this observation is new. Now Theorem 2.8 is a fundamental analytic fact about bounded kernels: it says that a bounded kernel is characterized up to equivalence by the quantities $s(F, \kappa)$, which are the natural analogues for a kernel of the moments of a random variable. When the first version of this paper was written, we thus had the following rather unsatisfactory situation: the only known proof of the analytic fact Theorem 2.8 was that given above, relying on the hard results of Borgs, Chayes, Lovász, Sós and Vesztergombi [15] about sequences of graphs. Fortunately, this situation has now been resolved: Borgs, Chayes and Lovász [12] have given a very clever direct proof of Theorem 2.8. In fact, they proved a little more.

Recall from (12) that $\kappa_1 \sim \kappa_2$ means that

$$\exists \kappa, \sigma_1, \sigma_2 \text{ such that } \kappa = \kappa_1^{(\sigma_1)} \text{ a.e. and } \kappa = \kappa_2^{(\sigma_2)} \text{ a.e.},$$

where κ is a kernel on $[0, 1]$ and σ_1 and σ_2 are measure-preserving maps from $[0, 1]$ to itself. Turning this ‘upside-down’, let us write $\kappa_1 \sim' \kappa_2$ if

$$\exists \kappa, \sigma_1, \sigma_2 \text{ such that } \kappa_1 = \kappa^{(\sigma_1)} \text{ a.e. and } \kappa_2 = \kappa^{(\sigma_2)} \text{ a.e.} \quad (16)$$

In (16), we require κ to be a kernel on $[0, 1]$; it makes no difference if we allow κ to be a kernel on an arbitrary standard probability space. Note that if $\kappa_1 \sim' \kappa_2$, then using the observation that $\kappa \sim \kappa^{(\sigma)}$ twice, we have $\kappa_1 \sim \kappa_2$.

Borgs, Chayes and Lovász [12] proved the following result.

Theorem 2.9. *For two bounded kernels κ_1, κ_2 , the following are equivalent. (a) $s(F, \kappa_1) = s(F, \kappa_2)$ for every finite graph F , (b) $\kappa_1 \sim \kappa_2$ and (c) $\kappa_1 \sim' \kappa_2$.*

The important implication is that if $s(F, \kappa_1) = s(F, \kappa_2)$ for all F , then $\kappa_1 \sim' \kappa_2$. As noted above, this trivially implies $\kappa_1 \sim \kappa_2$, which in turn easily implies $s(F, \kappa_1) = s(F, \kappa_2)$. The proof in [12] is direct, but somewhat technical.

As shown above, Theorem 2.9, which trivially implies Theorem 2.8, implies Theorem 2.4. This gives a proof of Theorem 2.4 that is very different from that given by Borgs, Chayes, Lovász, Sós and Vesztergombi [15].

Our aim in the rest of this paper is to investigate the extent to which the various results and observations above carry over to sparse graphs, graphs with n vertices and $o(n^2)$ edges. As we shall see, this gives rise to many difficult questions, so we shall present many more questions than answers.

3 Subgraph counts for sparse graphs

In this section we consider sparse graphs, where the number of edges is $o(n^2)$ as the number n of vertices goes to infinity. We shall assume throughout that we have at least $\omega(n)$ edges, i.e., that the average degree tends to infinity; often, we shall make much stronger assumptions. Given a function $p = p(n)$, one can adapt many of the notions of Section 2 to graphs with $\Theta(pn^2)$ edges. Indeed, let

$$s_p(F, G) = \frac{\text{emb}(F, G)}{p^{e(F)} n_{(|F|)}} = \text{aut}(F) \frac{X_F(G)}{p^{e(F)} X_F(K_n)},$$

noting that

$$s_p(F, G) = \frac{\text{emb}(F, G)}{\mathbb{E}(\text{emb}(F, G(n, p)))}.$$

Also, let

$$t_p(F, G) = \frac{\text{hom}(F, G)}{p^{e(F)} n^{|F|}}.$$

If $p = 1$, then we recover the definitions in Section 2. Furthermore, if $0 < p < 1$ is constant, then we can define a map s as before, but now s maps \mathcal{F} into the compact space $\prod_{F \in \mathcal{F}} [0, p^{-e(F)}]$, and everything proceeds as before. More generally, changing p by a constant factor will be irrelevant: just as we can use s_c for any c to study $G(n, 1/2)$, we may use s_p to study $G(n, p/2)$ or $G(n, 2p)$, say, for any $p = p(n)$.

From now on, we suppose that $p = p(n)$ is some given function of n , with $p(n) \rightarrow 0$ as $n \rightarrow \infty$. We wish to work in a compact space, so we shall assume that there are constants c_F , $F \in \mathcal{F}$, such that $s_p(F, G) \leq c_F$ for all graphs G we consider. Enumerating \mathcal{F} as $\{F_1, F_2, \dots\}$, we may thus define a map

$$s_p : \mathcal{F} \rightarrow X = \prod_{i=1}^{\infty} [0, c_{F_i}], \quad G \mapsto (s_p(F_i, G))_{i=1}^{\infty}, \quad (17)$$

and, using any metric d on X giving the product topology, an associated metric

$$d_{\text{sub}}(G_1, G_2) = d(s_p(G_1), s_p(G_2)). \quad (18)$$

We suppress the dependence on p in our notation for the metric to avoid clutter. As in the dense case, we can extend d_{sub} to bounded kernels κ , setting

$$d_{\text{sub}}(G, \kappa) = d(s_p(G), s(\kappa)) \quad \text{and} \quad d_{\text{sub}}(s(\kappa_1), s(\kappa_2)) = d(s(\kappa_1), s(\kappa_2))$$

for a graph G and bounded kernels κ , κ_1 and κ_2 . Here, for a kernel κ , $s(\kappa)$ is the vector with coordinates defined by (2).

Much of the time, we think of a sequence (G_n) of finite graphs. Throughout, we are only interested in sequences with $|G_n| \rightarrow \infty$. For notational convenience we always assume that $|G_n| = n$; this makes no difference to our conjectures and results. As usual, we need not assume that G_n is defined for every $n \in \mathbb{N}$, but only for an infinite subset of \mathbb{N} . In this setting, the assumption described above may be stated as follows.

Assumption 3.1 (bounded subgraph counts). For each fixed graph F , we have $\sup_n s_p(F, G_n) < \infty$.

In particular, if (G_n) satisfies Assumption 3.1 then, taking $F = K_2$, we see that $e(G_n) = O(pn^2)$, so our graphs are sparse. There is a stronger version of Assumption 3.1 that is perhaps even more natural:

Assumption 3.2 (exponentially bounded subgraph counts). There is a constant C such that, for each fixed F , we have $\limsup s_p(F, G_n) \leq C^{e(F)}$ as $n \rightarrow \infty$.

In this case, changing p by a constant factor, we may take $C = 1$ if we like. This is not always the most natural normalization, however. There is a reason for writing \limsup in Assumption 3.2: for any graph G_n with $|G_n| = n$ and n large, there will be some F with $s_p(F, G_n)$ very large. Indeed, G_n contains at least one embedding of itself, so $s_p(G_n, G_n) \geq 1/(n!p^{e(G_n)})$, which typically grows much faster than any constant to the power $e(G_n)$.

Turning to kernels, there is no longer any good reason to restrict our kernels to take values in $[0, 1]$: in the dense case, the maximum possible ‘local density’ of edges is 1. Here, if we normalize so that G_n has $pn^2/2$ edges, say, local densities larger than p are certainly possible. We shall thus consider general kernels, i.e., symmetric measurable functions from $[0, 1]^2$ to $[0, \infty)$, rather than only standard kernels. We define $s(F, \kappa)$ as before, using (2); in general, $s(F, \kappa)$ may be infinite, but we shall always assume it is finite for the graphs F and kernels κ we consider.

Although we allow unbounded kernels in general, it may be that they give rise to difficulties (as they do in the general (very) sparse inhomogeneous model of Bollobás, Janson and Riordan [8]). Assumption 3.2 corresponds to the limiting kernel (if it exists) being bounded, as shown by Lemma 3.5 below.

Our main conjecture states that, if p is large enough, then, under Assumption 3.2, the equivalent of Theorem 2.1 holds.

Conjecture 3.3. *Let $p = p(n) = n^{-o(1)}$, and let $C > 0$ be constant. Suppose that (G_n) is a sequence of graphs with $|G_n| = n$ such that, for every F , $s_p(F, G_n)$ converges to some constant $0 \leq c_F \leq C^{e(F)}$. Then there is a bounded kernel κ such that $c_F = s(F, \kappa)$ for every F .*

As noted above, without loss of generality we may take $C = 1$. As we shall observe later, it is very easy to see that if $s_p(K_2, G_n) \rightarrow 0$ and $s_p(F, G_n)$ is bounded for every F , then $s_p(F, G_n) \rightarrow 0$ for every F . Thus we may assume that $s_p(K_2, G_n)$ is bounded away from zero, and we may normalize in a different way by assuming that $s_p(K_2, G_n) = 1$, i.e. that $e(G_n) = p \binom{n}{2}$.

Assumption 3.2 is trivially stronger than Assumption 3.1. Thus, if (G_n) satisfies Assumption 3.2, then the sequence $s_p(G_n)$ defined by (17) lives in a compact product space, and has a convergent subsequence. Hence there are real numbers $c_F \geq 0$, $F \in \mathcal{F}$, and a subsequence (G_{n_i}) with $s_p(F, G_{n_i}) \rightarrow c_F$ for every F , to which Conjecture 3.3 applies. Conjecture 3.3 is thus a statement about the possible limit points of the sequences $s_p(G_n)$.

It may well be that the restriction to bounded kernels is not necessary.

Conjecture 3.4. *Let $p = p(n) = n^{-o(1)}$, and let (G_n) be a sequence of graphs with $|G_n| = n$ such that, for every F , we have $s_p(F, G_n) \rightarrow c_F$ for some $0 \leq c_F < \infty$. Then there is a kernel κ with $c_F = s(F, \kappa)$ for every F .*

We have stated the above conjectures under the assumption that $p = n^{-o(1)}$; we shall call this the *almost dense* case. The reason for this assumption is discussed further below. Let us note that, in the almost dense case, for each fixed F with k vertices, the denominator in the formula $\text{emb}(F, G_n)/(p^{e(F)}n_{(k)})$ for $s_p(F, G_n)$ is asymptotically $p^{e(F)}n^k$, which is $n^{k-o(1)}$. Since there are at most n^{k-1} non-injective homomorphisms from F to G_n , it follows that $t_p(F, G_n) \sim s_p(F, G_n)$ as $n \rightarrow \infty$, so it makes no difference whether we consider s_p or t_p . In general, this is not true: for example, considering homomorphisms which map all t vertices on one side of $K_{t,t}$ into a single vertex, we see that in any graph G_n with $pn^2/2$ edges there are at least $n(np)^t = (n^{2t}p^{t^2})/(np^t)^{t-1}$ non-injective embeddings of $K_{t,t}$. If np^t is bounded, then this is comparable to (or

larger than) the denominator in the definition of $t_p(K_{t,t}, G_n)$, and it follows that $t_p(K_{t,t}, G_n) - s_p(K_{t,t}, G_n)$ is bounded away from zero. Thus, for $t_p(K_{t,t}, G_n) \sim s_p(K_{t,t}, G_n)$ to hold with both quantities bounded, we need $np^t \rightarrow \infty$. This condition holds for every t only in the almost dense case $p = n^{-o(1)}$.

3.1 Bounded and unbounded kernels

The following simple observation illuminates the relationship between Conjectures 3.3 and 3.4.

Lemma 3.5. *Let $\kappa : [0, 1]^2 \rightarrow [0, \infty)$ be a kernel, and $C \geq 0$ a constant. Then we have $s(F, \kappa) \leq C^{e(F)}$ for every F if and only if $\kappa \leq C$ holds almost everywhere.*

Proof. The result is trivial if $C = 0$. Otherwise, rescaling, we may assume that $C = 1$. If $\kappa \leq 1$ almost everywhere, then $s(F, \kappa) \leq s(F, 1) = 1$ for every F . We may thus suppose that $\kappa > 1$ on a set of positive measure. It follows that there is some $\eta > 0$ such that $\kappa > (1 + \eta)^2$ on a set A of positive measure. Applying the Lebesgue Density Theorem to A , there is some $\varepsilon > 0$ and some rectangle $R = [a, a + \varepsilon] \times [b, b + \varepsilon] \subset [0, 1]^2$ such that $\mu(A \cap R) \geq \mu(R)/(1 + \eta)$. Thus, the average value of κ on the set R is at least $1 + \eta$. Let κ' be the kernel taking the value $1 + \eta$ on R and 0 elsewhere. Standard arguments from convexity show that, for each t ,

$$s(K_{t,t}, \kappa) \geq s(K_{t,t}, \kappa') = \varepsilon^{2t}(1 + \eta)^{t^2}.$$

Taking t large enough, we find an $F = K_{t,t}$ for which $s(F, \kappa) > 1$. \square

Lemma 3.5 shows that a kernel κ is bounded if and only if the counts $s(F, \kappa)$ grow at most exponentially in $e(F)$. It also shows that, in Conjecture 3.3, we need only consider kernels $\kappa : [0, 1]^2 \rightarrow [0, C]$.

Let us say that a kernel has *finite moments* if $s(F, \kappa) < \infty$ for all F . There are unbounded kernels with finite moments: the simplest way to construct such an example is to consider the ‘rank 1’ case, where $\kappa(x, y) = f(x)f(y)$ for some $f : [0, 1] \rightarrow [0, \infty)$. Indeed, let f be any function from $[0, 1]$ to $[0, \infty)$ with $\mathbb{E}(f^k) = \int_0^1 f(x)^k dx$ bounded for every k ; for example, let $f(x) = \log(1/x)$ for $x > 0$. Set $\kappa(x, y) = f(x)f(y)$. If F is a graph on $\{1, 2, \dots, k\}$ in which vertex i has degree d_i , then

$$\begin{aligned} s(F, \kappa) &= \int_{[0,1]^k} \prod_{ij \in E(F)} f(x_i)f(x_j) \prod_{i=1}^k dx_i \\ &= \int_{[0,1]^k} \prod_{i=1}^k f(x_i)^{d_i} \prod_{i=1}^k dx_i = \prod_{i=1}^k \mathbb{E}(f^{d_i}) < \infty. \end{aligned}$$

The calculation above shows that a rank one kernel $\kappa(x, y) = f(x)f(y)$ has finite moments if and only if $\|f\|_p < \infty$ for every $p > 1$, and hence if and only if

$\|\kappa\|_p < \infty$ for every $p > 1$. It is tempting to think that this holds in general. In one direction, for any kernel κ and any graph F on $\{1, 2, \dots, k\}$, we may write

$$s(F, \kappa) = \int_{[0,1]^k} \prod_{ij \in E(F)} \kappa_{ij}(x_1, \dots, x_d) \prod_{i=1}^k dx_i,$$

where $\kappa_{ij}(x_1, \dots, x_d) = \kappa(x_i, x_j)$. Thus, by Hölder's inequality,

$$s(F, \kappa) = \int \prod_{ij \in E(F)} \kappa_{ij} \leq \prod_{ij \in E(F)} \|\kappa_{ij}\|_{e(F)} = \prod_{ij \in E(F)} \|\kappa\|_{e(F)}.$$

Hence, if $\|\kappa\|_p < \infty$ for every $p > 1$, then $s(F, \kappa) < \infty$ for every F . The reverse implication does not hold, however, as shown by the following example.

Example 3.6. A kernel with finite moments but infinite 2-norm. Let us define a sequence of independent random kernels $\kappa_0, \kappa_1, \kappa_2, \dots$, as follows. For $r \geq 0$, let \mathcal{P}_r be the partition of $[0, 1]$ into 2^{2^r} equal intervals, and let \mathcal{P}_r^2 be the corresponding partition of $[0, 1]^2$: divide $[0, 1]^2$ into 4^{2^r} squares in the obvious way, and take as one part of \mathcal{P}_r^2 the union of a square and its reflection in the line $x = y$ (which may be the same square). Our kernel κ_r will be constant on each element of \mathcal{P}_r^2 , taking the value 2^r with probability 2^{-2^r} and 0 otherwise, with the values on different parts independent. Note that κ_0 is simply the constant kernel with value 1.

Let $\kappa(x, y) = \sum_{r=0}^{\infty} \kappa_r(x, y)$. It is easy to see that with probability 1 the sum converges almost everywhere (for example, recalling that μ denotes Lebesgue measure, use the fact that $\mathbb{E}\mu\{\kappa_r > 0\} = 2^{-2^r}$ to deduce that, with probability 1, $\mu\{\exists s > r : \kappa_s > 0\}$ tends to 0 as $r \rightarrow \infty$). Also, for large r , $\|\kappa_r\|_2^2$ is concentrated around its mean of $(2^r)^2 2^{-2^r} = 1$. Hence, with probability 1 we have $\|\kappa_r\|_2^2 \geq 0.99$ for infinitely many r . Using $(a + b)^2 \geq a^2 + b^2$ for $a, b \geq 0$, it follows that $\|\kappa\|_2^2$ is infinite with probability 1; in particular, κ does not have all p -norms finite.

Turning to the finite moments property, let F be any fixed graph, with t vertices. Since $\kappa \geq \kappa_0 = 1$, we have $s(F, \kappa) \leq s(K_t, \kappa)$, so we may assume without loss of generality that $F = K_t$. Since κ is random, $s(K_t, \kappa)$ is a random variable. We may write its expectation as

$$\mathbb{E}_{\kappa} \mathbb{E}_{\mathbf{x}} \prod_{i < j} \kappa(x_i, x_j) = \mathbb{E}_{\mathbf{x}} \mathbb{E}_{\kappa} \prod_{i < j} \kappa(x_i, x_j),$$

where \mathbb{E}_{κ} denotes expectation over the random choice of κ , and $\mathbb{E}_{\mathbf{x}}$ over the random choice of (x_1, \dots, x_t) , a sequence of t iid uniform elements of $[0, 1]$. Let us fix \mathbf{x} for the moment, assuming as we may that $x_i \neq x_j$ for $i \neq j$. Let ℓ be the largest r such that some pair x_i, x_j lie in the same part of \mathcal{P}_r , so $0 \leq \ell < \infty$. Let $\sigma = \sum_{r \leq \ell} \kappa_r$ and $\tau = \sum_{r > \ell} \kappa_r$, so $\kappa = \sigma + \tau$. For $r > \ell$, the $\binom{t}{2}$ pairs (x_i, x_j) , $i < j$, all lie in different parts of \mathcal{P}_r^2 , so the values of κ_r on these pairs are independent. Since different κ_r are independent, it follows that the values

of τ on the pairs are also independent. Now $\|\sigma\|_\infty \leq \sum_{r=0}^\ell \|\kappa_r\|_\infty = 2^{\ell+1} - 1$. Thus,

$$\mathbb{E}_\kappa \prod_{i < j} \kappa(x_i, x_j) \leq \mathbb{E}_\kappa \prod_{i < j} (2^{\ell+1} + \tau(x_i, x_j)) = \prod_{i < j} (2^{\ell+1} + \mathbb{E}_\kappa \tau(x_i, x_j)).$$

For any x and y we have $\mathbb{E}_\kappa \kappa_r(x, y) = 2^r 2^{-2r} = 2^{-r}$, from which it follows that $\mathbb{E}_\kappa \tau(x, y) \leq 2$, and hence, very crudely, that

$$\mathbb{E}_\kappa \prod_{i < j} \kappa(x_i, x_j) \leq \prod_{i < j} (2^{\ell+1} + 2) \leq 2^{2\ell t^2}.$$

It remains to take the expectation over \mathbf{x} . Since $\mathbb{P}(\ell = r) \leq \binom{t}{2} 2^{-2^r}$, we find that

$$\mathbb{E} s(F, \kappa) \leq \sum_{r=0}^\infty \binom{t}{2} 2^{-2^r} 2^{2rt^2} < \infty,$$

noting that for any fixed t the 2^{-2^r} term dominates. It follows that with probability 1 we have $s(F, \kappa) < \infty$ for every F , giving a kernel with finite moments but with $\|\kappa\|_2$ infinite. A simple modification, taking the probability that κ_r takes the value 2^r on a given square to be $2^{-(1+\varepsilon)r}$ rather than 2^{-2r} gives, for each $\varepsilon > 0$, an example with $\|\kappa\|_{1+\varepsilon}$ infinite.

3.2 Non-uniform random graphs

As in the dense case, there is a key connection between convergence of the counts $s_p(F, G_n)$ and random graphs. Given a kernel κ , let $G_p(n, \kappa)$ be the random graph on $[n]$ obtained as follows: first choose x_1, \dots, x_n independently and uniformly from $[0, 1]$. Then, conditional on this choice, join each pair $\{i, j\}$ of vertices independently, with probability $\min\{p\kappa(x_i, x_j), 1\}$. If $p\kappa$ is bounded by 1, then $G_p(n, \kappa)$ is simply $G(n, p\kappa)$; we write the parameter p as a subscript to emphasize that it is part of the overall normalization: we think of a sparse graph generated from the kernel κ , rather than a ‘sparse kernel’ $p\kappa$. If $p = 1/n$, then $G_p(n, \kappa)$ is a special case of the general sparse inhomogeneous model of Bollobás, Janson and Riordan [8].

Remark 3.7. In what follows, we shall consider many statements about the convergence of various sequences of random graphs. As usual in the theory of random graphs, the precise notion of convergence is not important: one thinks of ‘a random graph’ with certain asymptotic properties, although this makes no formal sense. Formally, it is most natural to work throughout with convergence in probability, but this would require us to consider ‘in probability’ versions of our various assumptions, for example the (exponentially) bounded counts assumptions 3.1 and 3.2. In fact, it is easy to check that in all cases considered here, the error probabilities decay fast enough to give almost sure convergence for any coupling of the relevant probability spaces. However, we shall not verify this explicitly, noting that one can in any case ensure almost sure convergence by passing to a suitable subsequence.

Lemma 3.8. *Let $p = p(n) = n^{-o(1)}$, and let κ be a kernel with $s(F, \kappa) < \infty$ for every F . Then $s_p(F, G_p(n, \kappa)) \xrightarrow{P} s(F, \kappa)$ for each fixed graph F , so $d_{\text{sub}}(G_p(n, \kappa), \kappa) \xrightarrow{P} 0$. In fact, the sequence $G_p(n, \kappa)$ converges almost surely to κ in the metric d_{sub} .*

Proof. It is very easy to check that, for every F , $s_p(F, G_p(n, \kappa))$ is concentrated around its mean $s(F, \kappa)$: indeed, the second moment of the number of copies of F can be written as a sum of terms $(1 + o(1))n_{(|H|)}p^{e(H)}s_p(H, \kappa)$, and the dominant term is the unique one with the largest power of n , where H is the disjoint union of two copies of F . (The $1 + o(1)$ correction is only needed if κ is unbounded, and appears due to the $\max\{1, \cdot\}$ in the edge probabilities.) This proves the first part of the result. Convergence in probability in d_{sub} follows since convergence in probability in a product topology is equivalent to convergence in probability of each coordinate. For the final statement, see Remark 3.7. \square

Lemma 3.8 implies that if κ has finite moments, then the sequence $G_n = G_p(n, \kappa)$ has bounded subgraph counts (i.e., satisfies Assumption 3.1) with probability 1. If κ is bounded, then G_n has exponentially bounded subgraph counts with probability 1.

Using Lemma 3.8, it is easy to see that we must allow unbounded kernels in Conjecture 3.4. Indeed, set $\kappa(x, y) = \log(1/x) \log(1/y)$ for $0 < x, y \leq 1$, say, and let $p(n) = 1/\log n$. Then the random graphs $G_p(n, \kappa)$ satisfy Assumption 3.1 with probability 1, and

$$s_p(F, G_p(n, \kappa)) \rightarrow s(F, \kappa) < \infty$$

holds with probability 1 for every F . Since κ is unbounded, by Lemma 3.5 there is no C with $s(F, \kappa) \leq C^{e(F)}$ for every F , so there is no bounded κ' with $s_p(F, G_p(n, \kappa)) \rightarrow s(F, \kappa')$ for every F .

Note that if p decreases too fast with n , then $s_p(F, G_p(n, \kappa))$ is no longer concentrated around its mean: for example, this is the case if $\mathbb{E} \text{emb}(F, G_p(n, \kappa))$ does not tend to infinity. This is the reason for the assumption $p = n^{-o(1)}$ in the various conjectures and results above: otherwise, there will be some F for which the expected number of embeddings does not tend to infinity. Note also that, for smaller p , when $s_p(F, \cdot)$ and $t_p(F, \cdot)$ are no longer asymptotically equal, the former is the more natural parameter: for a given F , the lower limit on p below which the corresponding parameter for $G_p(n, \kappa)$ is no longer close to $s(F, \kappa)$ is in general much smaller for $s_p(F, \cdot)$ than for $t_p(F, \cdot)$. It may well be, however, that the conjectures in this section (or perhaps just their proofs) fail when the relevant parameters $s_p(F, \cdot)$ and $t_p(F, \cdot)$ are no longer asymptotically equal.

3.3 Subgraph counts in the uniform case

Using convexity, it is very easy to check that the only possible kernel κ with $s(K_2, \kappa) = s(C_4, \kappa) = 1$ is the uniform kernel, with $\kappa = 1$ a.e. The following conjecture is thus a very special case of Conjecture 3.4.

Conjecture 3.9. *Let $p = p(n) = n^{-o(1)}$, and let (G_n) be a sequence of graphs with $|G_n| = n$, $e(G_n) = p\binom{n}{2}$, $s_p(C_4, G_n) \rightarrow 1$, and $\sup_n s_p(F, G_n) < \infty$ for each F . Then $s_p(F, G_n) \rightarrow 1$ for every F .*

Of course, there is a variant of Conjecture 3.9 where we replace Assumption 3.1 by Assumption 3.2, i.e., we demand that $\limsup_n s_p(F, G_n) \leq C^{e(F)}$ for some $C < \infty$. In this uniform context there is perhaps less reason to expect this to make a difference.

In the dense case, it is one of the basic results about quasi-random graphs that $s_p(K_2, G_n) \rightarrow 1$ and $s_p(C_4, G_n) \rightarrow 1$ imply $s_p(F, G_n) \rightarrow 1$ for every F , with no further assumptions; see Chung, Graham and Wilson [18]. In the sparse case, this result extends easily to certain graphs F ; here it turns out to be simpler to work with $t_p(F, G_n)$ rather than $s_p(F, G_n)$.

Lemma 3.10. *Let $p = p(n)$ with $pn^{1/2} \rightarrow \infty$, and let (G_n) be a sequence of graphs with $|G_n| = n$ such that $t_p(K_2, G_n) \rightarrow 1$ and $t_p(C_4, G_n) \rightarrow 1$. Then $t_p(C_k, G_n) \rightarrow 1$ for each $k \geq 5$.*

Proof. Suppressing the dependence on n , let A denote the adjacency matrix of G_n , and let $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$ be the eigenvalues of A . For $k \geq 3$ we have

$$\text{hom}(C_k, G_n) = \sum_{v_1, v_2, \dots, v_k \in V(G_n)} A_{v_1 v_2} A_{v_2 v_3} \cdots A_{v_k v_1} = \text{tr}(A^k) = \sum_{i=1}^n \lambda_i^k,$$

so

$$t_p(C_k, G_n) = n^{-k} p^{-k} \sum_{i=1}^n \lambda_i^k = \sum_{i=1}^n \mu_i^k, \quad (19)$$

where $\mu_i = \lambda_i/(np)$ is the i th normalized eigenvalue of G_n . In particular,

$$\sum_i \mu_i^4 \rightarrow 1. \quad (20)$$

The maximum eigenvalue of the adjacency matrix of any graph is at least the average degree, so

$$\mu_1 = (np)^{-1} \lambda_1 \geq (np)^{-1} (1 + o(1)) (n^2 p) / n = 1 + o(1).$$

From (20) it follows that $\mu_1 \sim 1$ and that $\sum_{i \geq 2} \mu_i^4 \rightarrow 0$. Hence $\mu_2 \leq 1$ and $\mu_n \geq -1$ if n is large enough, and then

$$\sum \mu_i^k = \mu_1^k + \sum_{i \geq 2} \mu_i^k \leq \mu_1^k + \max\{\mu_2^{k-4}, \mu_n^{k-4}\} \sum_{i \geq 2} \mu_i^4 \leq \mu_1^k + \sum_{i \geq 2} \mu_i^4 = 1 + o(1).$$

Using (19) again, the result follows. \square

Informally, when $pn^{1/2} \rightarrow \infty$, the parameters $s_p(C_k, G_n)$ and $t_p(C_k, G_n)$ are equivalent. More precisely, Lemma 3.10 implies the analogous statement with all occurrences of t_p replaced by s_p , but this requires a little work to show.

The restriction on p in Lemma 3.10 was not used in the proof. However, if G_n has average degree \bar{d} , then it contains at least $n\binom{\bar{d}}{2}$ pairs of adjacent edges. Thus, writing $N_{i,j}$ for the number of common neighbours of i and j , the sum of $N_{i,j}$ over ordered pairs $i \neq j$ is at least $2n\binom{\bar{d}}{2} = n\bar{d}(\bar{d}-1)$. Hence, the number of homomorphisms from C_4 to G_n with a given pair of opposite vertices mapped to distinct vertices is

$$\sum_{i \neq j} N_{i,j}^2 \geq \frac{1}{n(n-1)} \left(\sum N_{i,j} \right)^2 \geq \frac{n\bar{d}^2(\bar{d}-1)^2}{n-1}.$$

The number of homomorphisms with a given pair of opposite vertices mapped to the same vertex is simply the sum of the squares of the degrees in G_n , which is at least $n\bar{d}^2$. Thus,

$$\text{hom}(C_4, G_n) \geq \frac{n\bar{d}^2(\bar{d}-1)^2}{n-1} + n\bar{d}^2 \quad (21)$$

for any graph G_n with n vertices and average degree \bar{d} . With $\bar{d} \sim pn \rightarrow \infty$, this gives $\text{hom}(C_4, G_n) \geq (1+o(1))(n^4 p^4 + n^3 p^2)$, i.e., $t_p(C_4, G_n) \geq (1+o(1))(1+n^{-1}p^{-2})$. Consequently, $t_p(C_4, G_n) \sim 1$ implies $pn^{1/2} \rightarrow \infty$. When $pn^{1/2} \rightarrow \infty$, (21) reduces to the well-known fact that, in this case, $e(G_n) \sim p\binom{n}{2}$ implies that

$$t_p(C_4, G_n), s_p(C_4, G_n) \geq 1 - o(1).$$

In the dense case, Lemma 3.10 extends to triangles. Indeed, $\text{tr}(A^2)$ counts the number of walks of length 2 in G , which is just $2e(G)$. Thus

$$\sum \mu_i^2 = \frac{2e(G)}{n^2 p^2} \sim p^{-1}.$$

If p is bounded away from zero then it follows that $\sum_{i \geq 2} \mu_i^2$ is bounded as $n \rightarrow \infty$. Since $\sum_{i \geq 2} \mu_i^4 \rightarrow 0$, it follows by the Cauchy-Schwarz inequality that $\sum_{i \geq 2} \mu_i^3 \rightarrow 0$, and hence that $s_p(C_3, G_n) \rightarrow 1$.

To obtain a result for triangles in the sparse case by this method, one needs stronger assumptions. Defining p by $e(G) = n^2 p/2$, if we assume that $t_p(C_4, G_n) = 1 + o(p)$, then arguing as above we find that $\sum_{i \geq 2} \mu_i^4 = o(p)$ and $\sum_{i \geq 2} \mu_i^4 \leq p^{-1}$, so Cauchy-Schwarz does give $\sum_{i \geq 2} \mu_i^3 \rightarrow 0$. In general, many results for quasi-random graphs extend to the sparse case with similar modifications, where $o(1)$ error terms are replaced by suitable functions of p ; see, for example, the results of Thomason [37, 38] on (p, α) -jumbled graphs. Our aim here is different; we wish to assume only convergence in the relevant metric, making no assumption about the rate of convergence.

When $p \rightarrow 0$, the conditions of Lemma 3.10 do not guarantee the ‘right’ number of triangles, as our next two examples will show.

Example 3.11. Very sparse graphs with too few triangles. Throughout this example we assume that $p_1(n)$ and $p_2(n)$ are functions of n satisfying

$$p_2 = (1 - p_1^2)^{n-2} \quad (22)$$

and $p_1, p_2 = \Theta(\sqrt{\log n}/\sqrt{n})$. To be concrete, we may take $p_2 = \sqrt{\log n}/\sqrt{n}$, in which case the corresponding p_1 satisfies $p_1 \sim p_2/\sqrt{2}$. Suppressing the dependence on n , let G be the usual Erdős–Rényi random graph $G = G(n, p_1)$, and let H be the graph on the same vertex set $[n]$ in which vertices i and j are joined if and only if they do not have a common neighbour in G . From (22), each edge of H is present with probability p_2 ; note that the edges of H are *not* present independently of one another. For any set E of $r = O(1)$ possible edges of H , the edges of E are all present if and only if no vertex of G is joined to both ends of some edge in E . Considering each vertex of G separately, we see that the probability of this event is

$$(1 - rp_1^2 + O(p_1^3))^{n+O(1)} = e^{-rp_1^2n+O(np_1^3)} \sim ((1 - p_1^2)^{n-2})^r = p_2^r,$$

where the $O(1)$ correction in the first exponent is to account for vertices that are endpoints of one or more edges in E . In other words, the probability that a bounded number of edges is present in H is asymptotically the corresponding probability for $G(n, p_2)$.

For $E_1, E_2 \subset E(K_n)$, the event $E_2 \subset E(H)$ is a down-set in terms of G (it says that certain pairs of edges of G are not present), so $E_2 \subset E(H)$ and $E_1 \subset E(G)$ are negatively correlated. Hence, if $|E_2| = O(1)$, we have

$$\mathbb{P}(\{E_1 \subset E(G)\} \cap \{E_2 \subset E(H)\}) \leq (1 + o(1))p_1^{|E_1|}p_2^{|E_2|}. \quad (23)$$

Considering all ways of splitting a set E , it follows that $\mathbb{P}(E \subset G \cup H) \leq (1 + o(1))(p_1 + p_2)^{|E|}$, and hence that

$$\mathbb{E}(s_p(F, G \cup H)) \leq 1 + o(1) \quad (24)$$

for any fixed graph F , where $p = p_1 + p_2$.

Since G and H overlap in very few edges, and the numbers of edges of G and of H are concentrated, we have $s_p(K_2, G \cup H) \rightarrow 1$ almost surely. It follows that $s_p(C_4, G \cup H) \geq 1 - o(1)$ almost surely. Hence, from (24), $s_p(C_4, G \cup H) \xrightarrow{p} 1$, and it is not hard to deduce that $t_p(C_4, G \cup H) \xrightarrow{p} 1$.

On the other hand, there are by definition no triangles with two edges in G and one in H . Hence, from (23), the expectation of $\text{emb}(K_3, G \cup H)$ is at most

$$(1 + o(1))n^3(p_1^3 + 0 + 3p_1p_2^2 + p_2^3),$$

so $\mathbb{E}(s_p(K_3, G \cup H)) \leq (p^3 - 3p_1^2p_2)/p^3 + o(1)$. Since p_1, p_2 and p are all of the same order, this final fraction is strictly less than 1, and our construction gives almost surely a sequence $G_n = G \cup H$ with $s_p(K_2, G_n) \rightarrow 1$, $s_p(C_4, G_n) \rightarrow 1$ but $s_p(C_3, G_n) \not\rightarrow 1$. Since $\text{emb}(C_3, G) = \text{hom}(C_3, G)$ for any G , we have $t_p(C_3, G_n) \sim s_p(C_3, G_n) \not\rightarrow 1$. Choosing p_1 and p_2 satisfying (22) so that $p_2 \sim p_1/2$, we may achieve $s_p(C_3, G_n) \rightarrow 5/9$. Alternatively, choosing p_1 and p_2 suitably, we may find a sequence with $s_p(C_3, G_n) \not\rightarrow 1$ for any $p = p(n)$ satisfying $pn^{1/2} \rightarrow \infty$ and $p = O(\sqrt{\log n}/\sqrt{n})$.

Example 3.12. Very sparse graphs with no triangles. In the context of finding explicit constructions giving lower bounds on Ramsey numbers, Alon [1] constructed a sequence of graphs G_n defined only for certain n , with the following properties, where $d = d(n) \sim n^{2/3}/4$: the graph G_n is a d -regular Cayley graph, it is triangle free and (which is irrelevant here) the largest independent set has size $O(n^{2/3})$. In proving the last property, Alon shows that all eigenvalues other than $\lambda_1 = d$ are uniformly bounded by $O(n^{1/3})$. Setting $p = d/n$, so $t_p(K_2, G_n) = 1$, and writing μ_i for $\lambda_i/(np)$, as in the proof of Lemma 3.10, one thus has $\mu_1 = 1$ and $\mu_i = O(n^{-1/3})$ for $i \neq 1$, so from (19) it follows that $t_p(C_4, G_n) = 1 + O(n^{-1/3}) = 1 + o(1)$. This gives another example of a graph with almost the minimal number of C_4 s but too few (in this case no) triangles.

Example 3.13. Denser graphs with too few triangles. Let $n = mk$ where $m \rightarrow \infty$, and let $p = \sqrt{\log m}/\sqrt{m}$. Example 3.11 gives us a graph G' of order m with $t_p(K_2, G'), t_p(C_4, G') \sim 1$ and $t_p(K_3, G') \leq 0.9$, say, for all large enough m . Let G be the blow-up of G' obtained by replacing each vertex by k vertices. Since $t_p(F, \cdot)$ is unchanged by blow-ups, we have $t_p(K_2, G), t_p(C_4, G) \sim 1$ but $t_p(K_3, G) \leq 0.9$, from which $s_p(K_2, G), s_p(C_4, G) \sim 1$ and (for n large) $s_p(K_3, G) \leq 0.91$ follow immediately.

Although p has not changed, the number of vertices has. Seen as a function of n , we may choose $p = \sqrt{\log m}/\sqrt{m}$ for any m dividing n with $m \rightarrow \infty$. Exact divisibility is not essential. Either by using this fact, or by restricting to a subsequence, we see that any given function $p(n)$ can be realized up to a factor of $(1 + o(1))$, provided $p(n)/(\sqrt{\log n}/\sqrt{n}) \rightarrow \infty$ and $p(n) = o(1)$. Hence, we may construct graphs with the right number of C_4 s but too few triangles for any such function $p(n)$.

At first sight Example 3.13 seems to contradict Conjecture 3.9, but this is not the case. Indeed, for the graph G' that we blow up, (24) tells us that we do not have too many embeddings of any fixed F . However, while $s_p \sim t_p$ for $p = n^{-o(1)}$, the final p we consider, and while blowing up preserves t_p , G' is a very sparse graph: although it has the same absolute density as the final graph G , this density is much smaller than $|G'|^{-o(1)}$, since G' has many fewer vertices than G . It follows that the homomorphism counts in G' are *not* well behaved. In particular, G' contains around $m^4 p^3$ non-injective homomorphisms from $K_{2,3}$, which turns out to be much larger than the number $m^5 p^6$ of embeddings. It follows that G contains too many homomorphisms from, and thus embeddings of, $K_{2,3}$, i.e., that $s_p(K_{2,3}, G) \rightarrow \infty$.

Remark 3.14. Let us note in passing that the blowing-up argument above shows that replacing the assumption $p = n^{-o(1)}$ in Conjecture 3.9 (or Conjecture 3.3) with a stronger assumption such as $p(n) \geq 1/\log \log \log n$, say, makes no difference. Indeed, if the conjecture fails, and (G_n) is a counterexample, then blowing up G_n as above by replacing each vertex by $f(n)$ vertices for some rapidly growing $f(n)$ gives a counterexample for a different density function, where now the density goes to zero extremely slowly as a function of the number of vertices.

One possible approach to producing a counterexample to Conjecture 3.9 would be to consider *circulant graphs*, i.e., graphs on the vertex set $[n]$ in which whether or not ij is an edge depends only on $i - j$ modulo n . There is one circulant graph for each subset A of the integers modulo n satisfying $0 \notin A$ and $a \in A$ if and only if $-a \in A$. All our conjectures thus imply corresponding conjectures for subsets of \mathbb{Z}_n , the integers modulo n , in which the symmetry condition is not likely to be relevant. Most subgraph counts in the graph have a rather unnatural interpretation in terms of the corresponding sets; the exception is cycles, where the number of k -cycles in G corresponds to $(n$ times the) number of k -tuples in A^k summing to 0. There is a result corresponding to Lemma 3.10 for subsets of \mathbb{Z}_n , proved in the same way but using Fourier coefficients instead of eigenvalues. Unfortunately, Examples 3.11 and 3.13 also carry over to the set context, in a fairly straightforward way: instead of blowing up the graph, we replace each element of A by a block of consecutive integers. This shows that any result of the kind we want about subsets of \mathbb{Z}_n must involve conditions other than constraints on the number of tuples summing to 0.

In the sparse case, even when $p = n^{-o(1)}$, it is not true that $s_p(K_2, G_n) \rightarrow 1$ and $s_p(C_4, G_n) \rightarrow 1$ together imply $s_p(F, G_n) \rightarrow 1$ for every F . We have just seen one example, with $F = C_3$. There are also much simpler examples.

Example 3.15. Adding a dense part. Let $p = 1/\log n$, say, and let $m = m(n) = n/(\log n)^c$ where $c > 0$ is constant. (We ignore rounding to integers.) Let G' be any graph on $n - m$ vertices, and let G be the disjoint union of G' and a complete graph on m vertices. Since K_m contains roughly $m^{|F|}$ embeddings of any fixed F , we have

$$s_p(F, G) \sim s_p(F, G') + \frac{m^{|F|}}{p^{e(F)} n^{|F|}} = s_p(F, G') + (\log n)^{e(F) - c|F|}.$$

Taking $G' = G(n - m, p)$ and $c = 3/2$, say, we have $s_p(K_2, G) \sim s_p(K_2, G') \sim 1$, $s_p(C_4, G) \sim s_p(C_4, G') \sim 1$, but $s_p(K_4, G) \sim 1 + 1 = 2$. Note that $s_p(K_5, G) \rightarrow \infty$, so the assumptions of Conjecture 3.9 are not satisfied.

The above example is rather artificial: there are too many copies of K_4 (and of K_5), but these sit on a small number of vertices. However, the same effect can be achieved by taking the union on the same vertex set of $G(n, p)$ and a disjoint union of n/m copies of K_m . Also, we can use complete bipartite graphs instead of complete graphs.

Example 3.16. A blown-up random graph. Let $n = mk$, where $k = k(n)$ and $m = m(n)$ both tend to infinity. (As usual, we ignore divisibility issues, or consider a sequence $n_i \rightarrow \infty$.) Let G_1 be the random graph $G(m, p)$, where $p = p(n)$, and let $G = G_1^{(k)}$ be formed by replacing each vertex of G by an independent set of size k , and each edge by a k -by- k complete bipartite graph. The number of edges of G is $k^2 e(G_1)$, which is asymptotically $k^2 m^2 p/2 = n^2 p/2$, so $s_p(K_2, G) \rightarrow 1$ in probability and almost surely. Similarly, for any fixed graph F , each embedding of F into G_1 gives rise to $k^{|F|}$ embeddings

into G ; the expected number of embeddings arising in this way is essentially the expected number in $G(n, p)$, so whenever this expectation tends to infinity, such embeddings will contribute $1 + o(1)$ to $s_p(F, G)$.

There are other embeddings of F into G , however, where some distinct vertices of F are mapped to the same vertex in G_1 . For C_4 , we have roughly m^2pk^4 such embeddings within our complete bipartite graphs, and roughly $2m^3p^2k^4$ from embeddings involving three vertices of G_1 . Provided $mp^2 \rightarrow \infty$, we still have $s_p(C_4, G) \rightarrow 1$.

Fix an integer $t \geq 3$, and suppose now that $m = m(n)$ and $p = p(n)$ are chosen so that m and $k = n/m \rightarrow \infty$, and $mp^t \rightarrow c$ for some constant $0 < c < \infty$; for example, set $p = 1/\log n$, $m = c(\log n)^t$ and $k = c^{-1}n/(\log n)^t$. Note that $mp^2 \rightarrow \infty$. Then we have roughly $m^{2+t}k^{2+t}p^{2t}$ embeddings of $K_{2,t}$ into G coming from embeddings into G_1 . But we also have roughly $m^{1+t}k^{2+t}p^t$ embeddings into G coming from maps from $K_{2,t}$ into G_1 sending the two vertices on one side to the same vertex. It is easy to check that these two are the dominant terms (mapping the two vertices on one side to the same place we gain t factors of $1/p$ and lose one factor of m ; any other identifications gain fewer factors of $1/p$ per factor of m lost), and it follows that $s_p(K_{2,t}, G) \rightarrow 1 + 1/c$.

Taking a ‘typical’ sequence of random graphs constructed as above gives an example with $s_p(K_2, G_n) \rightarrow 1$, $s_p(C_4, G_n) \rightarrow 1$ (and indeed $s_p(K_{2,t'}, G_n) \rightarrow 1$ for $2 \leq t' < t$), but $s_p(K_{2,t}, G_n) \rightarrow 1 + 1/c \neq 1$. Once again, the assumptions of Conjecture 3.9 are not satisfied, this time because $s_p(K_{2,t+1}, G_n) \rightarrow \infty$.

We have seen from the examples above that if $p(n) \rightarrow 0$, then $s_p(K_2, G_n)$ and $s_p(C_4, G_n) \rightarrow 1$ do not themselves imply that $s_p(F, G_n) \rightarrow 1$ for every F . However, attempted counterexamples to Conjecture 3.9 seem to be doomed to failure by the additional assumption that $s_p(F, G_n)$ is bounded for every F . In the next section we shall see that we can make some progress towards proving Conjecture 3.9.

3.4 Partial results in the almost dense, uniform case

In the examples in the previous subsection, each vertex is in about the same number of copies of any fixed graph F , but there are relatively few ($o(n^2)$) pairs that are in too many copies of $K_{2,t}$, for example. It is easy to see that, under the assumptions of Conjecture 3.9, this cannot happen. In fact, we can make a much more general statement. For this it is convenient to work with homomorphism counts and $t_p(F, G_n)$ rather than embeddings and $s_p(F, G_n)$. As noted earlier, in the almost dense case that we consider in this subsection, i.e., when $p = n^{-o(1)}$, the quantities $t_p(F, G_n)$ and $s_p(F, G_n)$ differ by $o(1)$.

Let F be a fixed graph, and F' a subgraph of F . Without loss of generality, suppose that $V(F') = [\ell] \subset [k] = V(F)$. Then any homomorphism $\phi_F : F \rightarrow G_n$ restricts to a homomorphism $\phi_{F'} : F' \rightarrow G_n$. With $e(G_n) \sim p\binom{n}{2}$, we expect a typical $\phi_{F'}$ to have around $n^{k-\ell}p^{e(F)-e(F')}$ extensions. For each n , let us define a random variable $Z_n(F', F)$ as follows: let $\phi_{F'}$ be chosen uniformly at random from among all homomorphisms from F' into G (if there are any), and

let $Z_n(F', F)$ be the number of extensions of $\phi_{F'}$ divided by $n^{k-\ell}p^{e(F)-e(F')}$. (The reader may well prefer to picture copies of F' and F in G_n rather than homomorphisms. In fact, it is better to picture embeddings, i.e., labelled copies. There are essentially the same number of these as of homomorphisms.) Since $\text{hom}(F, G_n)$ is the sum over $\phi_{F'}$ of the number of extensions, we have

$$\text{hom}(F, G_n) = \text{hom}(F', G_n) \mathbb{E}(Z_n(F', F)) n^{k-\ell} p^{e(F)-e(F')},$$

and hence

$$t_p(F, G_n) = t_p(F', G_n) \mathbb{E}(Z_n(F', F)).$$

For $r \geq 2$, let rF/F' denote the graph formed by the union of r copies of F which all meet in the same subgraph F' , so rF/F' has $|F'| + r(|F| - |F'|)$ vertices and $e(F') + r(e(F) - e(F'))$ edges. A homomorphism from rF/F' to G_n consists of a homomorphism ϕ from F' to G_n together with r extensions of ϕ to homomorphisms from F to G , which may or may not be distinct. (They almost always will be.) Since we have normalized by the right powers of n and p , it follows that

$$t_p(rF/F', G_n) = t_p(F', G_n) \mathbb{E}(Z_n(F', F))^r. \quad (25)$$

Let $\mu_F = \mu_F(n) = n^{|F|} p^{e(F)}$, which is asymptotically equal to the expected number of homomorphisms from F into $G(n, p)$. Then, under the assumptions of any of Conjectures 3.3, 3.4 and 3.9, it is easy to see that for $F' \subset F$, any $o(\mu_{F'})$ copies of F' meet $o(\mu_F)$ copies of F . (Here ‘copies’ may be subgraphs of G_n , embeddings, or homomorphisms; it makes no difference.) Otherwise $t_p(2F/F', G_n)$ would not remain bounded. This rules out any construction of a potential counterexample similar to those above; it also shows that if $t_p(K_2, G_n) \rightarrow 0$ and Assumption 3.1 holds (i.e., (G_n) has bounded subgraph counts), then $t_p(F, G_n) \rightarrow 0$ for every F .

Conjecture 3.9 states that infinitely many conclusions (one for each F) hold under the same assumptions. We have already proved some of these conclusions, with $F = C_k$, $k \geq 5$. Our next aim is to prove a corresponding result for a much wider class of graphs. In doing so, the following observation will be useful.

Lemma 3.17. *Let $X_n \geq 0$ be a sequence of random variables with $\sup_n \mathbb{E}(X_n^k) < \infty$ for every $k \geq 1$. Then $\mathbb{E}(X_n^k) \rightarrow 1$ for every k if and only if $X_n \xrightarrow{P} 1$.*

Proof. For the forward implication we have $\mathbb{E}(X_n) \rightarrow 1$ and $\mathbb{E}(X_n^2) \rightarrow 1$; applying Chebyshev’s inequality it follows that $X_n \xrightarrow{P} 1$. The reverse implication is not much harder. Suppose that $X_n \xrightarrow{P} 1$, but that $\mathbb{E}(X_n^k) \not\rightarrow 1$ for some k . For any M , the variables $X_n^k 1_{X_n \leq M}$ are uniformly bounded and converge in probability to 1, so $\mathbb{E}(X_n^k 1_{X_n \leq M}) \rightarrow 1$. It follows that there is some $M(n) \rightarrow \infty$ such that $\mathbb{E}(X_n^k 1_{X_n \leq M(n)}) \rightarrow 1$. But then $\mathbb{E}(X_n^k 1_{X_n > M(n)}) \not\rightarrow 0$, so

$$\mathbb{E}(X_n^{k+1}) \geq \mathbb{E}(X_n^{k+1} 1_{X_n > M(n)}) \geq M(n) \mathbb{E}(X_n^k 1_{X_n > M(n)})$$

is unbounded, contradicting our assumptions. \square

Corollary 3.18. *Under the assumptions of Conjecture 3.9, if F' and F are fixed graphs with $F' \subset F$ and $t_p(F', G_n) \rightarrow 1$, then $Z_n(F', F) \xrightarrow{p} 1$ if and only if $t_p(rF/F', G_n) \rightarrow 1$ for every $r \geq 1$.*

Proof. Apply Lemma 3.17 to the random variable $Z_n(F', F)$, using (25) to evaluate its moments. \square

We shall say that the distribution of F is *flat over* that of F' in G_n , or simply that F is *flat over* F' , if $Z_n(F', F) \xrightarrow{p} 1$.

Lemma 3.19. *Under the assumptions of Conjecture 3.9 we have $s_p(K_{s,t}, G_n) \rightarrow 1$ for all $s, t \geq 1$. Moreover, $K_{1,s}$ is flat over E_s , where E_s is the empty subgraph of $K_{1,s}$ induced by the vertices in the second part.*

Proof. Let d_1, \dots, d_n denote the degrees of the vertices of G_n , and \bar{d} the average degree. Fix $s \geq 1$. By convexity, we have

$$\text{hom}(K_{1,s}, G_n) = \sum_{i=1}^n d_i^s \geq n\bar{d}^s,$$

which we can rewrite as $t_p(K_{1,s}, G_n) \geq t_p(K_2, G_n)^s$. Since $t_p(K_2, G_n) \rightarrow 1$ by assumption, this gives

$$t_p(K_{1,s}, G_n) \geq 1 + o(1). \quad (26)$$

Specializing to $s = 2$ for the moment, let $Z_n = Z_n(E_2, K_{1,2})$ be the random variable describing the distribution of the number of common neighbours of a random pair of vertices of G_n . For any empty graph E_k we have $t_p(E_k, G_n) = 1$. Hence, from (25) and (26),

$$\mathbb{E}(Z_n) = t_p(K_{1,2}, G_n) \geq 1 + o(1).$$

On the other hand, since $tK_{1,2}/E_2 = K_{2,t}$,

$$\mathbb{E}(Z_n^2) = t_p(K_{2,2}, G_n) = t_p(C_4, G_n) \rightarrow 1.$$

Since $\mathbb{E}(Z_n^2) \geq \mathbb{E}(Z_n)^2$, it follows that $\mathbb{E}(Z_n) \rightarrow 1$ and (by Lemma 3.17) that $Z_n \xrightarrow{p} 1$. In other words, $K_{1,2}$ is flat over pairs of vertices. By Corollary 3.18 it then follows that $t_p(K_{2,t}, G_n) \rightarrow 1$ for every t .

Returning to general s , let $W_n = Z_n(E_s, K_{1,s})$. From (26) we have $\mathbb{E}(W_n) = t_p(K_{1,s}, G_n) \geq 1 + o(1)$. But we have just shown that $\mathbb{E}(W_n^2) = t_p(K_{2,s}, G_n) \rightarrow 1$, so $W_n \xrightarrow{p} 1$, i.e., $K_{1,s}$ is flat over E_s . Applying Corollary 3.18 again we thus have $t_p(K_{s,t}, G_n) \rightarrow 1$ for every t , as required. \square

Theorem 3.20. *Let F be any fixed graph with girth at least 4, and let $F' \neq F$ be any induced subgraph of F . Under the assumptions of Conjecture 3.9, F is flat over F' . Furthermore, $s_p(F, G_n), t_p(F, G_n) \rightarrow 1$ as $n \rightarrow \infty$.*

Proof. Note first that the definition of $Z_n(F', F)$ makes perfect sense when F' is the empty ‘graph’ with no vertices; there is one homomorphism from F' to G_n , and $Z_n(F', F)$ is constant and takes the value $t_p(F, G_n)$. Hence, F is flat over the empty subgraph means exactly that $t_p(F, G_n) \rightarrow 1$. Since $p = n^{-o(1)}$, we have $s_p(F, G_n) \sim t_p(F, G_n)$, so it suffices to prove the first statement.

We prove the first statement of the theorem by induction on $|F|$. If $|F| = 1$, there is nothing to prove. Suppose then that F and F' are given, with $|F| \geq 2$, and that the result holds for all smaller F .

Suppose first that $F' = F - v$ for some vertex v of F . Let E_s denote the subgraph of F' induced by the neighbours of v , noting that E_s has no edges, as F is triangle free. Set $X_n = Z_n(E_s, F')$ and $Y_n = Z_n(E_s, K_{1,s})$. Note that these random variables are defined on the same probability space: the elements of this space are simply s -tuples of vertices of G_n . If $F' = E_s$, then F' is trivially flat over E_s . If not, then F' is flat over E_s by the induction hypothesis. Hence, in either case, $\mathbb{E}(X_n^k) \rightarrow 1$ for every k . By the last part of Lemma 3.19, $K_{1,s}$ is flat over E_s , so $\mathbb{E}(Y_n^k) \rightarrow 1$ for every k . It follows that $\mathbb{E}((X_n - 1)^k) \rightarrow 0$ and $\mathbb{E}((Y_n - 1)^k) \rightarrow 0$ for all $k \geq 1$. Hence, by the Cauchy–Schwarz inequality,

$$\mathbb{E}((X_n - 1)^k (Y_n - 1)^\ell) \leq \sqrt{\mathbb{E}((X_n - 1)^{2k}) \mathbb{E}((Y_n - 1)^{2\ell})} \rightarrow 0$$

for all $k, \ell \geq 0$ with $k + \ell > 0$. Writing $\mathbb{E}(X_n^k Y_n^\ell) = \mathbb{E}((X_n - 1 + 1)^k (Y_n - 1 + 1)^\ell)$ as 1 plus a sum of terms $\mathbb{E}((X_n - 1)^{k'} (Y_n - 1)^{\ell'})$, $k', \ell' \geq 0$, $k' + \ell' > 0$, it follows that $\mathbb{E}(X_n^k Y_n^\ell) \rightarrow 1$ for any $k, \ell \geq 0$.

Any homomorphism $\phi_{F'}$ from F' into G_n is the extension of a unique homomorphism ϕ_{E_s} from E_s into G_n . Furthermore, to extend $\phi_{F'}$ to F we must choose for the image of v a common neighbour of the vertices in the image of ϕ_{E_s} . Hence, the value of $Z_n = Z_n(F', F)$ on $\phi_{F'}$ is simply the value of Y_n on ϕ_{E_s} . Choosing $\phi_{F'}$ uniformly at random, to obtain the correct distribution for Z_n , the probability of obtaining a particular restriction ϕ_{E_s} is proportional to the number of extensions of ϕ_{E_s} to F' , i.e., to X_n . Thus the distribution of Z_n is that of Y_n ‘size biased’ by X_n . In particular,

$$\mathbb{E}(Z_n^k) = \frac{\mathbb{E}(X_n Y_n^k)}{\mathbb{E}(X_n)} \sim 1/1 = 1.$$

Taking $k = 1, 2$, it follows that $Z_n \xrightarrow{P} 1$, i.e., that F is flat over F' , as required.

It remains to handle the case $|F| - |F'| \geq 2$. In this case, we can find an induced subgraph $F'' = F - v$ of F with $F' \subset F'' \subset F$. Note that $t_p(F', G_n), t_p(F'', G_n) \sim 1$ by induction, that F'' is flat over F' by induction, and that F is flat over F'' by the case treated above. In particular, we certainly have

$$t_p(F, G_n) = t_p(F'', G_n) \mathbb{E}(Z_n(F'', F)) \sim 1.$$

Fix $\varepsilon > 0$. Let us call a copy of F'' (more precisely, a homomorphism from F'' into G_n) *bad* if it has fewer than $(1 - \varepsilon)\mu_F/\mu_{F''}$ extensions to copies of F . Since F is flat over F'' and $t_p(F'', G_n) \sim 1$, there are fewer than $\varepsilon^2 \mu_{F''}$ bad copies of

F'' if n is large enough. Since each copy of F'' extends a unique copy of F' , it follows that at most $\varepsilon\mu_{F'}$ copies of F' have more than $\varepsilon\mu_{F''}/\mu_{F'}$ extensions to bad copies of F'' .

Let \mathcal{B}_1 denote the set of copies of F' that have more than $\varepsilon\mu_{F''}/\mu_{F'}$ extensions to bad copies of F'' , so $|\mathcal{B}_1| \leq \varepsilon\mu_{F'}$ if n is large. Let \mathcal{B}_2 denote the set of copies of F' that have fewer than $(1-\varepsilon)\mu_{F''}/\mu_{F'}$ extensions to copies of F'' . Since F'' is flat over F' , we have $|\mathcal{B}_2| \leq \varepsilon\mu_{F'}$ if n is large enough, which we assume from now on. If ϕ is a copy of F' not in $\mathcal{B}_1 \cup \mathcal{B}_2$, then ϕ has at least $(1-2\varepsilon)\mu_{F''}/\mu_{F'}$ extensions to good copies of F'' , which in turn have at least $(1-\varepsilon)\mu_F/\mu_{F''}$ extensions to copies of F , so the value of $Z_n(F', F)$ on ϕ is at least $(1-2\varepsilon)(1-\varepsilon)$. Since there are $(1+o(1))\mu_{F'}$ copies of F' in total, the proportion of these copies in $\mathcal{B}_1 \cup \mathcal{B}_2$ is at most $\varepsilon + o(1)$. Since $\varepsilon > 0$ was arbitrary, it follows that the negative part of $Z_n(F', F) - 1$ tends to zero in probability. Since $\mathbb{E}(Z_n(F', F)) = t_p(F, G_n)/t_p(F', G_n) \rightarrow 1$, it follows that $Z_n(F', F) \xrightarrow{p} 1$, i.e., that F is flat over F' . \square

The reader may find many of the arguments above familiar from the dense case; for example, the proof for $K_{2,t}$ is an absolutely standard convexity argument. The key point is that many arguments for the dense case do not carry over. In particular, we have shown that almost all, i.e., all but $o(n^2)$, pairs of vertices have about the right number of common neighbours. In the dense case, it follows immediately that almost all (all but $o(pn^2) = o(n^2)$) edges are in the right number of triangles, and hence that $t_p(K_3, G_n) \rightarrow 1$. Similarly, the proof above shows that any F is flat over all its subgraphs in the dense case, without restriction to girth at least 4. In the sparse case, there are only $o(n^2)$ edges, and there seems to be no simple way to rule out the possibility that a large fraction, or even *all*, of the pairs of vertices corresponding to edges fall in the $o(n^2)$ set with too few common neighbours. Nevertheless, we conjecture that this cannot happen. The simplest graph for which we cannot prove the conclusion of Conjecture 3.9 is the triangle.

Conjecture 3.21. *Under the conditions of Conjecture 3.9 we have $s_p(K_3, G_n) \rightarrow 1$.*

In fact, we do not even have a proof that G_n must contain at least *one* triangle for n large enough!

3.5 Extensions to lower densities.

Let us return to the study of general subgraphs F , rather than simply triangles. If true, the various conjectures above may extend to smaller values of p , but one must be careful. Firstly, s_p and t_p no longer coincide, as noted above. One should work with s_p , because these quantities behave in the right way for $G_p(n, \kappa)$, while t_p does not. A simple modification of the proof of Lemma 3.19, considering the distribution of the number of common neighbours of a set of s *distinct* vertices, shows that if $np^s \rightarrow \infty$, then $s_p(K_2, G_n) \rightarrow 1$, $s_p(C_4, G_n) \rightarrow 1$ and $s_p(K_{s,t+1}, G_n)$ bounded together imply $s_p(K_{s,t}, G_n) \rightarrow 1$. Taking $p = n^{-\alpha}$,

with $0 < \alpha < 1/2$ constant, there is no corresponding result for t_p , even with $s = 2$. Indeed, if $t_p(K_2, G_n) = 1$, then there are at least $n^{t+1}p^t$ homomorphisms from $K_{2,t}$ into G_n mapping the two vertices in the smaller class to the same vertex. It follows that $t_p(K_{2,t}, G_n)$ will be unbounded for any $t > 1/\alpha$.

Secondly, even working with s_p rather than t_p , we cannot in general hope to conclude in the analogue of Conjecture 3.4 that $s_p(F, G_n) \rightarrow s(F, \kappa)$ for *all* fixed graphs F . For example, set $p = n^{-1/2}$ and consider the *polarity graphs* G_n of Erdős and Rényi [22], defined (for suitable n) by taking as vertices the points of the projective plane over $GF(q)$, q a prime power, and joining $x = (x_0, x_1, x_2)$ and $y = (y_0, y_1, y_2)$ if and only if $x_0y_0 + x_1y_1 + x_2y_2 = 0$ in $GF(q)$. These graphs satisfy $e(G_n) \sim n^{3/2}/2 = pn^2/2$ but contain no C_4 s, and thus satisfy $s_p(K_2, G_n) \rightarrow 1$ and $s_p(C_4, G_n) = 0$. Since $s(C_4, \kappa) \geq s(K_2, \kappa)^4$ for any κ , we cannot have $s_p(F, G_n) \rightarrow s(F, \kappa)$ for $F = K_2$ and for $F = C_4$ in this case. More generally, whenever $pn^{1/2} \not\rightarrow \infty$, then there are graphs G_n with pn^2 edges but too few C_4 s, so we should only consider the counts $s_p(C_4, G_n)$ if $pn^{1/2} \rightarrow \infty$. This problem is not unique to C_4 , so it seems that to extend our conjectures for $p = n^{-o(1)}$ to sparser graphs, we should modify them to refer only to a certain set of ‘admissible’ subgraphs F , depending on the function $p = p(n)$.

In fact, we should only consider subgraphs F for which the expected number $\mu_F \sim n^{|F|}p^{e(F)}$ of embeddings of F into $G(n, p)$ is much larger than the number $(1 + o(1))n^2p/2$ of edges, at least if $pn^{1/2} \rightarrow \infty$. To see this, first suppose that $n^{|F|}p^{e(F)} \sim An^2p$, for some constant $0 < A < \infty$. Form a graph G' from $G = G(n, p)$ by adding $\varepsilon n^2p/(2e(F))$ copies F_1, F_2, \dots of F , chosen uniformly at random from all subgraphs of K_n isomorphic to F . After deleting the small number of duplicate edges, we have added around $\varepsilon n^2p/2$ edges, so $s_p(K_2, G') \sim 1 + \varepsilon$. It is easy to check that the number of C_4 s in G' containing two or more edges from one single F_i is negligible and thus, considering C_4 s formed from all combinations of edges from $G(n, p)$ and from different F_i , that $s_p(C_4, G') \sim (1 + \varepsilon)^4$ whp. Hence, the appropriate limiting kernel is the constant kernel $\kappa = 1 + \varepsilon$. Copies of F itself containing at most one edge from each F_i contribute $(1 + \varepsilon)^{e(F)}$ to $s_p(F, G')$, but there are $\Theta(n^{|F|}p^{e(F)})$ extra copies of F , namely the F_i themselves. It follows that $s_p(F, G_n) \not\rightarrow 1$. If $n^{|F|}p^{e(F)} = o(n^2p)$, then the argument is much simpler: adding a few copies of F to $G(n, p)$ does not change the number of edges or C_4 s significantly, but does change the number of copies of F .

We can go somewhat further: the construction in Example 3.11 shows that for C_3 to be admissible, the expected number of C_3 s per edge should be larger than $\log n$. A similar construction can be carried out for any fixed F , and shows that, at least for suitable balanced F , we should require $n^{|F|}p^{e(F)}/(np^2 \log n) \rightarrow \infty$ for F to be admissible. In general, for F to be admissible, we need all induced subgraphs F' of F to be admissible; otherwise, the distribution of copies of F over F' cannot be flat as we expect in the uniform case.

Returning to triangles, in the light of the comments above, perhaps the strongest conceivable extension of Conjecture 3.21 to smaller p would be that if $p = p(n) = \omega(\sqrt{\log n}/\sqrt{n})$, and $s_p(K_2, G_n) \rightarrow 1$, $s_p(C_4, G_n) \rightarrow 1$, and $\sup_n s_p(K_{2,t}, G_n) < \infty$ for each t , then $s_p(C_3, G_n) \rightarrow 1$. However, it may well

be that the graphs constructed by Alon [1] mentioned in Example 3.12 have $s_p(K_{2,t}, G_n) \rightarrow 1$ for each t . (This may also be true of Kim's random construction [30] giving his famous lower bound on the Ramsey numbers $R(3, t)$.) If so, blowing these graphs up as in Example 3.13 would show that even in the almost dense case, controlling the $K_{2,t}$ counts is not enough, so one should control (at least) the $K_{s,t}$ counts for some larger s . Returning to much sparser graphs, we then have to limit ourselves to $p = p(n)$ for which $K_{s,t}$ is admissible, suggesting the following conjecture.

Conjecture 3.22. *There are constants $s \geq 2$ and $a > 0$ such that, if $p = p(n) = \omega((\log n)^a n^{-1/s})$ and G_n is a sequence of graphs with $|G_n| = n$, $s_p(K_2, G_n) \rightarrow 1$, $s_p(C_4, G_n) \rightarrow 1$, and $\sup_n s_p(K_{s,t}, G_n) < \infty$ for each t , then $s_p(C_3, G_n) \rightarrow 1$.*

It may be that if the conjecture holds for a given s , it holds with $c = 1/s$. It may also be that one needs to control the counts for $K_{s,t}$ and at the same time to consider p larger than n^{-b} for some $b < 1/s$.

There is a potential pitfall in handling subgraph counts when p is smaller than $n^{-o(1)}$: in proving that $s_p(F, G_n) \rightarrow 1$ for various graphs F above, we made use of the assumption that $s_p(F', G_n)$ is bounded for other graphs F' . In particular, with $F = K_{2,t}$, we used this assumption for $F' = K_{2,t+1}$. It may be that F' is admissible whenever F is (as is likely in this case: $K_{2,t}$ should be admissible as soon as C_4 is), but perhaps not. In the latter case we may be forced to work with a larger admissible set for which we impose the hypothesis of Conjecture 3.3 (or Conjecture 3.4), and a smaller set for which we obtain the conclusion. In any case, the (smaller) admissible set should have the following property: if \mathcal{F}_α denotes the set of admissible graphs when $p = n^{-\alpha}$, $\alpha > 0$, then the sets \mathcal{F}_α should increase as α decreases, and their union should contain all finite graphs. We shall return to this question in Section 5, in particular in Subsections 5.3 and 5.4, where we prove results that are steps towards (non-uniform) versions of the various conjectures in this section.

4 Szemerédi's Lemma and the cut metric

In the next section we shall discuss the relationship between the cut and count metrics. As in the dense case, a key tool in the study of the cut metric is some variant of Szemerédi's Lemma [36]: this will be discussed in this section. Unlike in the dense case, we need an assumption on the graphs we consider to make this useful; roughly speaking, our assumption is that no subgraph of G_n containing a constant fraction of the vertices has density more than a constant factor larger than it should have. Several of the usual proofs of Szemerédi's Lemma extend easily to the sparse case under this assumption; this was noted independently by Kohayakawa and Rödl; see [32]. (The much earlier Theorem 2 of Kohayakawa [31] is slightly different.)

Throughout this section, $p = p(n)$ with $p = o(1)$ and $np \rightarrow \infty$. (Often, $n^2p \rightarrow \infty$ is enough in the proofs, but see Remark 4.4.) As before, (G_n) always

denotes a sequence of graphs with $|G_n| = n$, which need not be defined for all n , but only for some infinite set.

For disjoint sets A, B of vertices of a graph $G = G_n$ with n vertices, we write $e_G(A, B)$ for the number of edges of G joining A to B , and

$$d_p(A, B) = \frac{e_G(A, B)}{p|A||B|} \quad (27)$$

for the normalized density of G between A and B . It is convenient to extend this definition to sets A and B that need not be disjoint: in this case, we write $e_G(A, B)$ for the number of ordered pairs (i, j) with $i \in A, j \in B$ and $ij \in E(G)$; we then define $d_p(A, B)$ as above. Note that $e_G(A, A) = 2e(G[A])$. We shall make the following assumption:

Assumption 4.1 (bounded density). There is a constant C and a function $n_0(\varepsilon)$ such that, for every $\varepsilon > 0$ and $n \geq n_0(\varepsilon)$, and any $A, B \subset V(G_n)$ with $|A|, |B| \geq \varepsilon n$, we have $d_p(A, B) \leq C + \varepsilon$.

It suffices to impose this assumption only when $A = B$, replacing C by $C/2$ and ε by $\varepsilon/2$. Indeed, if $|A|, |B| \geq \varepsilon n$, $n \geq n_0(\varepsilon)$, and $d_p(A, B) > C + \varepsilon$ then, by averaging, we may find $A' \subset A$ and $B' \subset B$ with $|A'| = |B'| = \lceil \varepsilon n \rceil$ such that $d_p(A', B') > C + \varepsilon$. Then $e_G(A' \cup B', A' \cup B') \geq 2e_G(A', B') > 2(C + \varepsilon)|A'|^2 \geq (C/2 + \varepsilon/2)|A' \cup B'|^2$.

The condition above may be written more compactly as follows:

$$\forall \varepsilon > 0 : \limsup_{n \rightarrow \infty} \max\{d_p(A, B) : A, B \subset V(G_n), |A|, |B| \geq \varepsilon n\} \leq C. \quad (28)$$

Note that we shall often assume that (28) holds for a particular value of C : in this case, we say that (G_n) has *density bounded by C* . This is the reason for including the final $+\varepsilon$ in Assumption 4.1.

It will be convenient to phrase the proof of Szemerédi's Lemma in terms of kernels. In this sparse setting, the way in which we associate a kernel to a graph is different from in the dense case. Indeed, our aim is that the random graph $G(n, p)$ should approximate the constant kernel taking value 1. For this reason, to a graph G with n vertices $1, 2, \dots, n$ we associate the kernel κ_G taking the value $1/p$ on each square $((i-1)/n, i/n] \times ((j-1)/n, j/n]$ whenever $ij \in E(G)$, and zero elsewhere. This association will often be implicit: for example, given a graph G and a kernel κ , we write $d_{\text{cut}}(G, \kappa)$ for $d_{\text{cut}}(\kappa_G, \kappa)$.

The following observation shows the importance of bounded density. In the proof, and throughout this section, given a subset A of the vertices of a graph G , we shall often abuse notation by also writing A for the corresponding subset of $[0, 1]$.

Lemma 4.2. *Let $p = p(n)$ be any function of n , let $\kappa : [0, 1]^2 \rightarrow [0, C]$ be a kernel, and let (G_n) be a sequence of graphs with $|G_n| = n$ and $d_{\text{cut}}(G_n, \kappa) \rightarrow 0$. Then (G_n) has density bounded by C .*

Proof. Suppose that (G_n) does not have density bounded by C . Then there is an $\varepsilon > 0$ such that, for infinitely many n , there are sets $A_n, B_n \subset V(G_n)$ with $|A_n|, |B_n| \geq \varepsilon n$ and $d_p(A_n, B_n) \geq C + \varepsilon$. Identifying A_n and B_n with subsets of $[0, 1]$, and writing μ for Lebesgue measure, we have

$$\int_{A_n \times B_n} \kappa_{G_n} = d_p(A_n, B_n) \mu(A_n) \mu(B_n) \geq (C + \varepsilon) \mu(A_n) \mu(B_n).$$

Since κ is bounded by C , it follows that

$$\left| \int_{A_n \times B_n} \kappa_{G_n} - \kappa^{(\tau)} \right| \geq \varepsilon \mu(A_n) \mu(B_n) \geq \varepsilon^3$$

for any rearrangement $\kappa^{(\tau)}$ of κ , which contradicts $d_{\text{cut}}(G_n, \kappa) \rightarrow 0$. \square

4.1 Weakly regular partitions

If G is a graph with vertex set $\{1, 2, \dots, n\}$, and $\Pi = (P_1, \dots, P_k)$ is a partition of $V(G)$, then we write G/Π for the kernel on $[0, 1]^2$ taking the value $d_p(P_a, P_b)$ on the union of the squares $((i-1)/n, i/n] \times (j-1)/n, j/n]$, $i \in P_a, j \in P_b$. We say that a partition Π of a graph G is *weakly* (ε, p) -*regular* if $\|\kappa_G - G/\Pi\|_{\text{cut}} \leq \varepsilon$. Note that the normalizing function p comes in via the definition of the kernels κ_G and G/Π .

For a kernel κ , the definitions are similar: for $A, B \subset [0, 1]$ we write $\kappa(A, B)$ for the integral of κ over $A \times B$, and

$$d(A, B) = d_\kappa(A, B) = \frac{\kappa(A, B)}{\mu(A)\mu(B)}$$

for the average value of κ on $A \times B$. Then $d_p(A, B)$, defined using G , is exactly $d(A, B)$, defined using κ_G , so the kernel G/Π is obtained from κ_G by replacing the value at each point by the average over the relevant rectangle $P_a \times P_b$. For κ a kernel and Π a partition of $[0, 1]^2$, we define κ/Π similarly. The partition Π is *weakly* ε -*regular* with respect to κ if $\|\kappa - \kappa/\Pi\|_{\text{cut}} \leq \varepsilon$.

The next lemma is a sparse equivalent of (a version of) the Frieze–Kannan ‘weak’ form of Szemerédi’s Lemma from [23]. As with many proofs of the various forms of Szemerédi’s Lemma, the proof of the dense result is not hard to adapt to the sparse setting: the only additional complication is that one must make sure that the parts of the partition remain large enough so that we can make use of the bounded density assumption. In the following lemma, $p = p(n)$ is any normalizing function with $pn^2 \rightarrow \infty$. In principle, the various constants depend on the choice of p , but this is not the case if we impose an explicit lower bound on $p(n)$, such as the harmless bound $p \geq n^{3/2}$.

Lemma 4.3. *Let $p = p(n)$ be any function with $0 < p \leq 1$ and $pn^2 \rightarrow \infty$. Let $\varepsilon > 0$, $C > 0$ and $k \geq 1$ be given. There exist constants n_0 , K and $\eta > 0$, all depending on ε , C and k , such that, if G_n is any graph with $n \geq n_0$ vertices such that*

$$d_p(A, B) \leq C \text{ whenever } |A|, |B| \geq \eta n, \quad (29)$$

and Π is any partition of $V(G)$ into k parts P_1, \dots, P_k with sizes as equal as possible, then there is a weakly (ε, p) -regular partition Π' of $V(G_n)$ into K parts that refines Π .

Proof. Reducing ε if necessary, we may assume that $\varepsilon \leq C$, say. We assume without comment that n is ‘large enough’ whenever this is needed.

Let $\Pi_0 = \Pi$. We shall inductively define a sequence Π_t of partitions of $V(G)$ into $k_t = 2^t k$ parts, stopping either when we reach some Π_t that is weakly $(\varepsilon/2, p)$ -regular, or when $t \geq T = \lceil 16C^2/\varepsilon^2 \rceil + 1$. Every part of Π_t will have size at least $\gamma^t n/(2k)$, where $\gamma = \varepsilon/(100C) \leq 1/100$. Note that Π_0 satisfies this condition.

Set $\eta = \gamma^T/(2k)$, and let n_0 be a large constant to be chosen later. We shall write κ_t for the kernel G/Π_t , noting that, since all parts of Π_t have size at least ηn , the kernel κ_t is bounded by C .

Given Π_t as above, suppose that Π_t is not weakly $(\varepsilon/2, p)$ -regular. Then there is a cut $[0, 1] = A \cup A^c$ exhibiting this, i.e., a set $A \subset [0, 1]$ for which $|\kappa_G(A, A^c) - \kappa_t(A, A^c)| \geq \varepsilon/2$. Since both κ_G and κ_t correspond to weighted graphs on $V(G) = \{1, 2, \dots, n\}$, we may choose the cut A to correspond to a subset of $V(G)$: among all ‘worst’ cuts, there is a cut of this form.

Our aim is to modify A slightly to obtain a set B (which we may think of as a subset of $V(G)$ or as a subset of $[0, 1]$) and then take two parts $P_i \cap B$ and $P_i \cap B^c$ of Π_{t+1} for each part of Π_t ; in doing so, we must ensure that neither of these parts is too small. We modify the set A to obtain B in k_t stages, one for each part P_i . At each stage, we move a set S of at most $\gamma|P_i| \geq \eta n$ vertices from A to A^c or vice versa, to ensure that both B and B^c meet P_i in at least $\gamma|P_i|$ vertices. Since κ_t is bounded by C , this changes the value of the cut $\kappa_t(A, A^c)$ by at most $2C\gamma|P_i|/n$.

From (29), the set S meets at most $Cpn\gamma|P_i|$ edges of G : to see this, apply (29) to S and $V(G)$ if $|S| \geq \eta n$, and to S' and $V(G)$ otherwise, for any $S' \supset S$ with $\lceil \eta n \rceil$ vertices. Hence, the value of the cut $\kappa_G(A, A^c)$ changes by at most $2C\gamma|P_i|/n$ when we move our set S from one side of the cut to the other. After all these changes, we have

$$|\kappa_t(A, A^c) - \kappa_t(B, B^c)|, |\kappa_G(A, A^c) - \kappa_G(B, B^c)| \leq 2C\gamma \leq \varepsilon/8.$$

It follows that

$$|\kappa_G(B, B^c) - \kappa_t(B, B^c)| \geq \varepsilon/4. \quad (30)$$

Let Π_{t+1} be the partition obtained by intersecting each part of Π_t with B and B^c , noting that Π_{t+1} has all the required properties. Set $\kappa_{t+1} = G/\Pi_{t+1}$, noting that $\kappa_{t+1}(B, B^c) = \kappa_G(B, B^c)$, since Π_{t+1} refines the partition (B, B^c) . From (30) it thus follows that

$$\|\kappa_{t+1} - \kappa_t\|_1 \geq \|\kappa_{t+1} - \kappa_t\|_{\text{cut}} \geq \varepsilon/4,$$

with the final inequality witnessed by the cut (B, B^c) . Hence, $\|\kappa_{t+1} - \kappa_t\|_2^2 \geq \|\kappa_{t+1} - \kappa_t\|_1^2 \geq \varepsilon^2/16$. Since κ_t may be obtained from κ_{t+1} by averaging over rectangles, κ_t and $\kappa_{t+1} - \kappa_t$ are orthogonal: for any two parts P_i, P_j of Π_t , the

kernel κ_t is constant on $P_i \times P_j$. Also, $\int_{P_i \times P_j} \kappa_{t+1} = \int_{P_i \times P_j} \kappa_G = \int_{P_i \times P_j} \kappa_t$. Thus $\int_{P_i \times P_j} \kappa_t(\kappa_{t+1} - \kappa_t) = 0$. Summing over i and j it follows that $\int \kappa_t(\kappa_{t+1} - \kappa_t) = 0$. Thus,

$$\|\kappa_{t+1}\|_2^2 = \|\kappa_t\|_2^2 + \|\kappa_{t+1} - \kappa_t\|_2^2 \geq \|\kappa_t\|_2^2 + \varepsilon^2/16.$$

It follows by induction that $\|\kappa_t\|_2^2 \geq t\varepsilon^2/16$ as long as our construction continues. But, as noted above, κ_t is bounded by C , so our construction must stop after at most $16C^2/\varepsilon^2$ steps. Since this number is smaller than T , we must stop at a weakly $(\varepsilon/2, p)$ -regular partition.

To complete the proof we modify the final partition Π_t slightly. Set $K = k\lceil\gamma^{-T}\rceil$, and note that, since $t \leq T - 1$, each part of Π_t has size at least $\gamma^{-1}n/K$. First, adjust the parts slightly so that the size of each is of the form $a\lceil n/K \rceil + b\lfloor n/K \rfloor$, $a, b \in \mathbb{Z}^+$, replacing the kernel κ_t by a new kernel κ' corresponding to the altered partition Π' . Arguing as above, $\|\kappa_t - \kappa'\|_{\text{cut}} \leq 2C\gamma \leq \varepsilon/4$, so, by the triangle inequality and weak $(\varepsilon/2, p)$ -regularity of Π_t , we have

$$\|\kappa' - \kappa_G\|_{\text{cut}} \leq \|\kappa_t - \kappa_G\|_{\text{cut}} + \|\kappa_t - \kappa'\|_{\text{cut}} \leq \varepsilon/2 + \varepsilon/4 = 3\varepsilon/4.$$

Finally, we split each part randomly into parts of sizes exactly $\lfloor n/K \rfloor$ and $\lceil n/K \rceil$, obtaining a partition Π'' into K parts whose sizes are as equal as possible. We write κ'' for the corresponding kernel. Since Π' has $O(1)$ parts, and we have $\Theta(pn^2)$ edges between any two parts with density at least $\varepsilon/100$, say, it follows from Chernoff's inequality that if n is large enough, which we enforce by choosing n_0 suitably, then with probability at least 0.99 the density $d_p(A, B)$ between every pair (A, B) of new parts A and B coming from parts P_i and P_j of Π' with $d_p(P_i, P_j) \geq \varepsilon/100$ is $d_p(P_i, P_j)(1 + o(1))$. Since the densities $d_p(P_i, P_j)$ are uniformly bounded by C , it follows that with probability at least 0.99 we have $\|\kappa'' - \kappa'\|_1 \leq \varepsilon/100$. But then

$$\|\kappa'' - \kappa_G\|_{\text{cut}} \leq \|\kappa'' - \kappa'\|_{\text{cut}} + \|\kappa' - \kappa_G\|_{\text{cut}} \leq \|\kappa'' - \kappa'\|_1 + \|\kappa' - \kappa_G\|_{\text{cut}} \leq \varepsilon/100 + \varepsilon/2,$$

so our final partition Π'' is indeed weakly (ε, p) -regular. \square

If for any reason we want a weakly (ε, p) -regular partition into a particular number K of parts (which must be a multiple of the number in the original partition if we are refining a given partition), the proof above gives such a partition for any large enough K , indeed, for any $K \geq k\lceil\gamma^{-T}\rceil$. Of course, n_0 then depends on K .

Remark 4.4. The proof of Lemma 4.3 works even if p is very small, say of order $1/n$. However, this is of no help – it is impossible for Assumption 4.1 (the sequence version of (29)) to be satisfied in this range, except in the trivial case where $e(G_n) = o(pn^2)$ (so p is not the appropriate normalizing function). Indeed, passing to a subsequence where $e(G_n)/(pn^2)$ is bounded away from zero, picking any εpn^2 edges of G_n , and putting one endpoint of each edge into A and the other into B , we find sets A, B with $|A|, |B| \leq \varepsilon pn^2$ but $e(A, B) \geq \varepsilon pn^2$, which gives $d_p(A, B) \geq 1/(\varepsilon p^2 n^2)$, which tends to infinity as $\varepsilon \rightarrow 0$.

4.2 Strongly regular partitions

Usually, when working with the cut metric, weak ε -regularity turns out to be just as good as the usual stronger ε -regularity. In the dense case, this is true also when considering subgraph counts. However, for the subgraph counts we consider in the next section, it turns out that we do in fact need the usual form of ε -regularity.

As usual, a pair (A, B) of (not necessarily disjoint) subsets of $V(G)$ is an (ε, p) -regular pair if $|d_p(A', B') - d_p(A, B)| \leq \varepsilon$ whenever $A' \subset A$ and $B' \subset B$ satisfy $|A'| \geq \varepsilon|A|$ and $|B'| \geq \varepsilon|B|$. A partition $\Pi = (P_1, \dots, P_k)$ of $V(G)$ is (ε, p) -regular if the parts P_i each have size $\lceil n/k \rceil$ or $\lfloor n/k \rfloor$, and all but at most $\varepsilon \binom{k}{2}$ of the unordered pairs $\{P_i, P_j\}$, $i \neq j$, are (ε, p) -regular. The definition (now simply of ε -regularity) for a kernel is similar, although here one partitions the interval $[0, 1]$ into parts with measure exactly $1/k$.

The following is (essentially) the sparse version of Szemerédi's Lemma observed by Kohayakawa and Rödl; see [32], where a closely related result is proved. For a proof, see also Gerke and Steger [24]. We shall include a proof here as we state the result in a slightly different way (which makes no real difference), and the use of kernels allows one to phrase the proof a little more simply than in [32] or [24].

Lemma 4.5. *Let $p = p(n)$ be any function with $0 < p \leq 1$ and $pn^2 \rightarrow \infty$. Let $\varepsilon > 0$, $C > 0$ and $k \geq 1$ be given. There exist constants n_0 , K and $\eta > 0$, all depending on ε , C and k , such that, if G_n is any graph with $n \geq n_0$ vertices such that*

$$d_p(A, B) \leq C \text{ whenever } |A|, |B| \geq \eta n, \quad (31)$$

and Π is any partition of $V(G)$ into k parts P_1, \dots, P_k with sizes as equal as possible, then there is an (ε, p) -regular partition Π' of $V(G_n)$ into at most K parts that refines Π .

Proof. Reducing ε and/or increasing C if necessary, we may suppose for convenience that $\varepsilon \leq 1$ and $C \geq 1$.

Set $\gamma = \varepsilon^3/(100C)$. This time we inductively define a sequence Π_t of partitions of $V(G)$ into k_t parts, where $\Pi_0 = \Pi$, $k_0 = k$, and $k_{t+1} = k_t \lceil k_t 2^{k_t} / \gamma \rceil$, stopping either when we reach some Π_t that is (ε, p) -regular, or when $t \geq T = \lceil 20C^2/\varepsilon^5 \rceil + 1$. The parts of each Π_t will have sizes as equal as possible. Note that Π_0 satisfies this condition.

Set $\eta = 1/(2k_T)$, and let n_0 be a large constant to be chosen later. We assume throughout that $n \geq n_0$. As before, we write κ_t for the kernel G/Π_t , noting that, since all parts of Π_t have size at least ηn , the kernel κ_t is bounded by C .

The key (standard) observation is the following. Let A and B be parts of Π_t , so κ_t is by definition constant on $A \times B$, and let $A' \subset A$ and $B' \subset B$. Let Π' be any partition refining Π such that each of A' and B' is a union of parts of Π' , and let $\kappa' = G/\Pi'$ be the corresponding kernel. Restricted to $A \times B$, the function κ' integrates to $d_p(A, B)\mu(A)\mu(B) = \int_{A \times B} \kappa_t$, since A

and B are unions of parts of κ' . Hence, κ_t and $\kappa' - \kappa_t$ are orthogonal on this set. Using the fact that A' and B' are unions of parts of κ' , we see that $\int_{A' \times B'} \kappa' = d_p(A', B')\mu(A')\mu(B')$, which differs from the integral of κ_t over the same set by $|d_p(A', B') - d_p(A, B)|\mu(A')\mu(B')$. It follows that $\|\kappa' - \kappa_t\|_2^2$ is at least $(d_p(A', B') - d_p(A, B))^2\mu(A')\mu(B')$, and hence, using orthogonality, that

$$\int_{A \times B} (\kappa')^2 \geq \int_{A \times B} \kappa_t^2 + (d_p(A', B') - d_p(A, B))^2\mu(A')\mu(B'). \quad (32)$$

Suppose then that Π_t is not (ε, p) -regular, and let A_1, \dots, A_{k_t} denote the parts of Π_t . Then there are at least $\varepsilon \binom{k_t}{2}$ pairs $\{A_i, A_j\}$ of parts of Π_t that are not (ε, p) -regular. For each, pick sets $A_{ij} \subset A_i$ and $A_{ji} \subset A_j$ witnessing this, i.e., with $|d_p(A_{ij}, A_{ji}) - d_p(A_i, A_j)| \geq \varepsilon$ and $|A_{ij}| \geq \varepsilon|A_i|$, $|A_{ji}| \geq \varepsilon|A_j|$. Let Π' be the partition whose parts are all atoms formed by the sets A_i and the sets A_{ij} taken together, so Π' refines Π_t , and each A_{ij} is a union of parts of Π' . We could estimate the L^2 -norm of G/Π' using (32), but this will not be useful if some parts of Π' are too small, so we first adjust the part sizes.

Define Π_{t+1} by dividing each A_i into k_{t+1}/k_t parts whose sizes are as equal as possible, so that each part of Π' differs from a union of parts of Π_{t+1} in at most n/k_{t+1} vertices: to do this, keep taking for a part of Π_{t+1} a subset of some part of Π' , until what is left of every part of Π' is too small. For each i , there are at most k_t sets A_{ij} inside A_i , so A_i is a union of at most 2^{k_t} parts of Π' . It follows that there is some union A'_{ij} of parts of Π_{t+1} with

$$|A_{ij} - A'_{ij}| \leq 2^{k_t}n/k_{t+1} \leq \gamma n/k_t^2.$$

Arguing as in the proof of Lemma 4.3, it follows from (31) that the symmetric difference S_{ij} of A_{ij} and A'_{ij} meets at most

$$Cpn|S_{ij}| \leq Cp\gamma n^2/k_t^2 \leq \varepsilon^3 p|A_i||A_j|/99$$

edges of G , if n is sufficiently large. Since $|S_{ij}| \leq \varepsilon^3|A_i|/100$, say, while $|A_{ij}| \geq \varepsilon|A_i|$ and $|A_{ji}| \geq \varepsilon|A_j|$, it follows crudely that

$$|d_p(A'_{ij}, A'_{ji}) - d_p(A_{ij}, A_{ji})| \leq \varepsilon/2,$$

which implies that

$$|d_p(A'_{ij}, A'_{ji}) - d_p(A_i, A_j)| \geq \varepsilon/2.$$

Now A'_{ij} and A'_{ji} are unions of parts of Π_{t+1} , and these sets have size at least $\varepsilon n/(2k_t)$. Hence, from (32),

$$\int_{A_i \times A_j} \kappa_{t+1}^2 \geq \int_{A_i \times A_j} \kappa_t^2 + \varepsilon^4/(16k_t^2)$$

for each of the at least $\varepsilon \binom{k_t}{2}$ irregular pairs $\{A_i, A_j\}$. Since $\int_{A_i \times A_j} \kappa_{t+1}^2 \geq \int_{A_i \times A_j} \kappa_t^2$ always holds, it follows that $\|\kappa_{t+1}\|_2^2 \geq \|\kappa_t\|_2^2 + \varepsilon^5/20$.

If the construction above does not stop before step T , then by induction we have $\|\kappa_t\|_2 \geq t\varepsilon^5/20$ for $0 \leq t \leq T$. But each κ_t is bounded by C , so $\|\kappa_T\|_2^2 \leq C^2$, giving a contradiction. Hence the construction does stop before step T , giving an (ε, p) -regular partition with $k_t \leq k_T$ parts. \square

Note that Lemma 4.5 implies (essentially) Lemma 4.3: it is easy to check that an (ε, p) -regular partition is, say, weakly $(10(C+1)\varepsilon, p)$ -regular, provided the parts are large enough for (31) to hold. However, one of course obtains much worse bounds on the number of parts using the stronger notion of regularity.

Remark 4.6. Let us illustrate once again the difference between the dense and sparse cases with a simple observation. Given a pair (A, B) of sets of vertices of a graph G , let $C_4(A, B)$ denote the number of homomorphisms from C_4 into the subgraph spanned by $A \cup B$ mapping a given pair of opposite vertices into A and the other pair into B . Standard convexity arguments show that $C_4(A, B) \geq d(A, B)^4 |A|^2 |B|^2$. The pair (A, B) is (ε, p) - C_4 -minimal if $C_4(A, B) \leq (d(A, B)^4 + \varepsilon p^4) |A|^2 |B|^2$. In the dense case (with $p = 1$) it is well known and very easy to check that ε -regularity and ε - C_4 -minimality are essentially equivalent: ε -regularity implies $f(\varepsilon)$ - C_4 -minimality, and ε - C_4 -minimality implies $g(\varepsilon)$ -regularity, for some $f(\varepsilon), g(\varepsilon)$ with $f(\varepsilon), g(\varepsilon) \rightarrow 0$ as $\varepsilon \rightarrow 0$.

Let $\varepsilon > 0$ and M be given. By counting C_4 s it is easy to see that there is a function $f(\varepsilon)$ with $f(\varepsilon) \rightarrow 0$ as $\varepsilon \rightarrow 0$ such that, if n is large enough and (A, B) is ε -regular with $|A| = |B| = n$, then we may partition A and B into sets A_1, \dots, A_M and B_1, \dots, B_M of almost equal sizes so that every pair (A_i, B_j) is $f(\varepsilon)$ -regular. Indeed, a random partition has this property with probability tending to 1, since by standard concentration results (for example, the Hoeffding–Azuma inequality), the edge densities and ‘ C_4 -densities’ of the pairs (A_i, B_j) are highly concentrated about the corresponding densities for (A, B) . It follows immediately that in the usual dense Szemerédi’s Lemma [36], we may specify in advance the number of parts K we would like our partition to have, provided (as in the weak case) that K is large enough given ε , and n large enough given ε and K .

In the sparse case, the fact about random partitioning above is presumably true, but the simple proof using C_4 -counts fails totally. It is still true that (ε, p) - C_4 -minimality implies $(f(\varepsilon), p)$ -regularity, but the reverse implication fails. Indeed, whenever $p = p(n) \rightarrow 0$, given any pair (A, B) , we may add a small dense (say complete bipartite) subgraph with too few edges to disturb regularity, but containing many more than $p^4 |A| |B|$ C_4 s.

4.3 Szemerédi’s lemma and convergence in the cut norm

We start with a consequence of Lemma 4.3 concerning the cut norm.

Corollary 4.7. *Let (G_n) be a sequence of graphs satisfying Assumption 4.1. Then there is a kernel $\kappa : [0, 1]^2 \rightarrow [0, C]$ and a subsequence (G_{n_i}) of (G_n) such that $d_{\text{cut}}(G_{n_i}, \kappa) \rightarrow 0$. Moreover, we may label the vertices of G_{n_i} with $1, 2, \dots, n_i$ so that $\|\kappa_{G_{n_i}} - \kappa\|_{\text{cut}} \rightarrow 0$.*

Proof. We shall only sketch the proof as the argument is exactly the same as that of Lovász and Szegedy [34] for the dense case. Note that given any $\eta > 0$ and $\varepsilon > 0$, our graphs G_n satisfy the assumption (29) of Lemma 4.3 with $C + \varepsilon$ in place of C whenever n is large enough.

First, let us apply Lemma 4.3 with $k = 1$ and $\varepsilon = \varepsilon_1 = 1/2$, say, to obtain a weakly (ε_1, p) -regular partition $\Pi_{n,1}$ of G_n into $k_1 = K$ parts, for all large enough n . We may relabel the vertices of each G_n so that the parts of $\Pi_{n,1}$ are all intervals. Each kernel $G_n/\Pi_{n,1}$ is characterized by a k_1 -by- k_1 density matrix, whose entries all lie in $[0, C + \varepsilon_1]$. (Indeed, if k_1 happens to divide n , then the kernel is exactly the kernel obtained from the matrix in the obvious way.) Since these matrices live in a compact set, $[0, C + \varepsilon_1]^{k_1^2}$, they have a convergent subsequence. Passing to the corresponding subsequence of G_n , we then have $G_n/\Pi_{n,1} \rightarrow \kappa_1$ pointwise almost everywhere, and hence in L^1 and in the cut norm. Since the partitions $\Pi_{n,1}$ are weakly (ε_1, p) -regular, we have $\|\kappa_{G_n} - G_n/\Pi_{n,1}\|_{\text{cut}} \leq \varepsilon_1$. Passing far enough along our subsequence, it follows that $\|\kappa_{G_n} - \kappa_1\|_{\text{cut}} \leq 2\varepsilon_1$.

Working within the subsequence defined above, apply Lemma 4.3 again with $\varepsilon = \varepsilon_2 = 1/4$, say, and $k = k_1$. For each n we find a partition $\Pi_{n,2}$ refining $\Pi_{n,1}$, with $k_2 = K(\varepsilon_1, C, k_1)$ parts. Relabelling vertices, we may assume that each part of each $\Pi_{n,2}$ is an interval. (Note that we only reorder the vertices within parts of $\Pi_{n,1}$.) As before, on a subsequence we have $G_n/\Pi_{n,2} \rightarrow \kappa_2$, for some kernel κ_2 constant on squares of side-length $1/k_2$. Since $\Pi_{n,2}$ refines $\Pi_{n,1}$ for each n , it follows that the value of κ_1 on each $1/k_1$ -by- $1/k_1$ square is exactly the average of κ_2 over this set; to see this, let $n \rightarrow \infty$.

Iterating, we find kernels $\kappa_1, \kappa_2, \dots$ each of which can be obtained by averaging the next one, and graphs G_{n_i} with $\|\kappa_{G_{n_i}} - \kappa_i\|_{\text{cut}} \leq 2\varepsilon_i = 2^{1-i}$, say. To complete the proof we simply observe that the sequence (κ_t) is a martingale on the state space $[0, 1]^2$. Since each κ_t is bounded by $C + \varepsilon_t \leq C + 1$, by the Martingale Convergence Theorem there is a kernel $\kappa : [0, 1]^2 \rightarrow [0, C]$ with $\kappa_t \rightarrow \kappa$ pointwise almost everywhere, and hence in L^1 and in the cut-norm. Then $\|\kappa_{G_{n_i}} - \kappa\|_{\text{cut}} \rightarrow 0$ as required. \square

The corollary above says that any (suitable) sequence of graphs has a subsequence converging to a kernel, and is a simple consequence of Szemerédi's Lemma and the Martingale Convergence Theorem. Together with Lemma 4.2, it shows that Assumption 4.1 is the correct assumption to impose on sequences of graphs when we seek limits that are bounded kernels $\kappa : [0, 1]^2 \rightarrow [0, C]$. Before turning to an application of Corollary 4.7, let us note an even simpler consequence of the Martingale Convergence Theorem.

Lemma 4.8. *Let κ be a bounded kernel, and for $k \geq 1$, let κ_k be the piecewise constant kernel obtained by dividing $[0, 1]^2$ into 2^{2k} squares of side 2^{-k} , and replacing κ by its average over each square. Then $\kappa_k \rightarrow \kappa$ pointwise almost everywhere and also in L^p for any p .*

Proof. The sequence κ_k is a bounded martingale on $[0, 1]^2$, so pointwise convergence is given by the Martingale Convergence Theorem. Since the sequence κ_k

is bounded by $\sup \kappa$, convergence in L^p follows by dominated convergence. \square

A consequence of Corollary 4.7 is that it allows us to compare the two different versions of the cut metric. Recall that for graphs G_1, G_2 , we defined $d_{\text{cut}}(G_1, G_2)$ by first passing to kernels taking the values 0 and $1/p$. If G_1 and G_2 have the same number of vertices, then there is a more natural definition of their cut-distance, $\widehat{d}_{\text{cut}}(G_1, G_2)$, defined in the same way but only allowing rearrangements that ‘map whole vertices to whole vertices’. As in the dense case, $d_{\text{cut}}(G_1, G_2)$ and $\widehat{d}_{\text{cut}}(G_1, G_2)$ are defined by (8) and (9), respectively; the difference between the sparse and dense cases is in the normalization of κ_{G_i} . Writing d_{cut}^1 and $\widehat{d}_{\text{cut}}^1$ for the metrics defined using $p = 1$, Borgs, Chayes, Lovász, Sós and Vesztegombi [15, Theorem 2.3] showed that these metrics are equivalent, proving that

$$d_{\text{cut}}^1(G_1, G_2) \leq \widehat{d}_{\text{cut}}^1(G_1, G_2) \leq 32d_{\text{cut}}^1(G_1, G_2)^{1/67}. \quad (33)$$

In fact, they proved (33) for edge-weighted graphs, as long as all edge weights lie in $[-1, 1]$. Unlike simple Lipschitz equivalence, which may also hold, this does not directly carry over to the sparse setting: we have $d_{\text{cut}} = p^{-1}d_{\text{cut}}^1$ and $\widehat{d}_{\text{cut}} = p^{-1}\widehat{d}_{\text{cut}}^1$, so (33) can be written as

$$d_{\text{cut}}(G_1, G_2) \leq \widehat{d}_{\text{cut}}(G_1, G_2) \leq 32d_{\text{cut}}(G_1, G_2)^{1/67}p^{-66/67},$$

which is of little if any use here. However, the equivalence of the two metrics in the sparse case is not too hard to deduce from (33), using Corollary 4.7.

Lemma 4.9. *For $i = 1, 2$, let $(G_n^{(i)})$ be a sequence of graphs satisfying the bounded density assumption 4.1. Then $d_{\text{cut}}(G_n^{(1)}, G_n^{(2)}) \rightarrow 0$ if and only if $\widehat{d}_{\text{cut}}(G_n^{(1)}, G_n^{(2)}) \rightarrow 0$.*

Proof. If $\widehat{d}_{\text{cut}}(G_n^{(1)}, G_n^{(2)}) \rightarrow 0$ then, since $d_{\text{cut}} \leq \widehat{d}_{\text{cut}}$, it follows trivially that $d_{\text{cut}}(G_n^{(1)}, G_n^{(2)}) \rightarrow 0$.

Suppose now that $d_{\text{cut}}(G_n^{(1)}, G_n^{(2)}) \rightarrow 0$; our aim is to show that $\widehat{d}_{\text{cut}}(G_n^{(1)}, G_n^{(2)}) \rightarrow 0$, so we may suppose that this is not the case. Hence, passing to a subsequence, we may assume that $\widehat{d}_{\text{cut}}(G_n^{(1)}, G_n^{(2)}) \geq \delta$ for some positive δ and all n in our subsequence.

Applying Corollary 4.7 twice, the second time to a suitable subsequence, we find kernels $\kappa_1, \kappa_2 : [0, 1]^2 \rightarrow [0, C]$, and subsequences of the sequences $(G_n^{(i)})$, defined for the same values of n , on which $\|\kappa_{G_n^{(i)}} - \kappa_i\|_{\text{cut}} \rightarrow 0$. Since $d_{\text{cut}}(G_n^{(1)}, G_n^{(2)}) \rightarrow 0$, it follows that $d_{\text{cut}}(\kappa_1, \kappa_2) = 0$.

For any $\varepsilon > 0$, by Lemma 4.8 we may find a K and kernels $\kappa'_1, \kappa'_2 : [0, 1]^2 \rightarrow [0, C]$ that are constant on squares of side $1/K$, with $\|\kappa'_i - \kappa_i\|_{\text{cut}} \leq \varepsilon$. Since the kernels κ'_i may be thought of as weighted graphs, it would appear that we have gone round in circles, but the point is that they are *dense* weighted graphs. Regarding the kernels κ'_1/C and κ'_2/C as weighted graphs with edge weights in $[0, 1]$, we have

$$d_{\text{cut}}(\kappa'_1/C, \kappa'_2/C) = d_{\text{cut}}(\kappa'_1, \kappa'_2)/C \leq 2\varepsilon/C,$$

so (33) gives

$$\widehat{d}_{\text{cut}}(\kappa'_1, \kappa'_2) = C\widehat{d}_{\text{cut}}(\kappa'_1/C, \kappa'_2/C) \leq 32C(2\varepsilon/C)^{1/67} = O(\varepsilon^{1/67}).$$

Hence, there is a rearrangement of κ'_1 preserving intervals that is close to κ'_2 in the cut norm. Ignoring divisibility, adapting this rearrangement to the graph $G_n^{(1)}$, n much larger than K , and using $\|\kappa_{G_n^{(i)}} - \kappa'_i\|_{\text{cut}} \leq \varepsilon + o(1)$, it follows that $\widehat{d}_{\text{cut}}(G_n^{(1)}, G_n^{(2)}) \leq O(\varepsilon) + O(\varepsilon^{1/67})$. Choosing ε small enough, the final bound is less than δ , contradicting our assumptions. \square

Corollary 4.7 shows that one property of the cut metric carries over to the sparse setting: for every suitable sequence (G_n) , i.e., any sequence satisfying the bounded density assumption 4.1, there is a kernel κ and a subsequence converging to κ in d_{cut} . In the other direction, as in the dense case, such a sequence is given by the natural random construction.

Lemma 4.10. *Let $p = p(n)$ satisfy $np \rightarrow \infty$, let $C > 0$ be constant, let $\kappa : [0, 1]^2 \rightarrow [0, C]$ be a bounded kernel, and let $G_n = G_p(n, \kappa)$. Then $d_{\text{cut}}(G_n, \kappa) \rightarrow 0$ almost surely. Also, the sequence (G_n) satisfies the bounded density assumption 4.1 with probability 1.*

Proof. The second statement is essentially immediate from Chernoff's inequality, constructing G_n as a subgraph of the Erdős-Rényi random graph $G(n, Cp)$; it also follows from the first statement and Lemma 4.2.

We now turn to the proof that $d_{\text{cut}}(G_n, \kappa) \rightarrow 0$. Recall that κ is of finite type if $[0, 1]$ may be partitioned into sets A_1, \dots, A_k so that κ is constant on each rectangle $A_i \times A_j$. We first suppose that κ is of finite type. Rearranging κ , and ignoring parts with measure zero, we may assume that each A_i is an interval with positive measure. Recall that $G_n = G_p(n, \kappa)$ is constructed by first choosing the 'types' x_1, \dots, x_n of the vertices independently and uniformly at random from $[0, 1]$. Let n_i denote the number of vertices of type i , noting that we have $n_i \sim \mu(A_i)n$ a.s. Let us adjust the intervals A_i slightly, replacing A_i by a set $A'_i (= A'_i(n))$ with measure n_i/n . Let $\kappa' = \kappa'(n)$ be the adjusted kernel, taking on $A'_i \times A'_j$ the value that κ takes on $A_i \times A_j$. Since, a.s., we adjust the length of each A_i by $o(1)$, the kernels κ' and κ differ on a set of measure $o(1)$. Since each is bounded, it follows that

$$\|\kappa - \kappa'\|_{\text{cut}} \leq \|\kappa - \kappa'\|_1 \rightarrow 0 \quad (34)$$

a.s., as $n \rightarrow \infty$.

Given x_1, \dots, x_n , let G' be the weighted graph in which each edge is present and has weight $w_{ij} = p\kappa(x_i, x_j)$. Then, relabelling the vertices so that those with $x_i = k$ correspond to the set A'_k , we see that $\kappa_{G'} = \kappa'$. The graph G_n may be constructed from G' by simply selecting each edge ij independently, with probability equal to its weight in G' . As noted earlier, for a kernel corresponding to a (weighted) graph, the cut norm (defined by (3)) is realized by a cut

corresponding to a partition of the vertex set, so

$$\|\kappa_{G_n} - \kappa'\|_{\text{cut}} = \|\kappa_{G_n} - \kappa_{G'}\|_{\text{cut}} = \max_{S \subset V(G_n)} \left| \frac{e_{G_n}(S, S^c) - \sum_{i \in S, j \in S^c} w_{ij}}{n^2 p} \right|.$$

Having conditioned on x_1, \dots, x_n , for each S the random variable $X = e_{G_n}(S, S^c)$ has mean exactly $\sum_{i \in S, j \in S^c} w_{ij}$. Furthermore, $\mathbb{E}(X) = O(n^2 p)$. Since X is a sum of independent indicator variables, it follows from (for example) the Chernoff bounds, that for any $\varepsilon > 0$ we have $\mathbb{P}(|X - \mathbb{E}(X)| \geq \varepsilon n^2 p) \leq \exp(-c_\varepsilon n^2 p)$ for some $c_\varepsilon > 0$. Since $n^2 p = \omega(n)$, this probability decays superexponentially. Since there are only 2^n sets S to consider, we see that $\mathbb{P}(\|\kappa_{G_n} - \kappa'\|_{\text{cut}} \geq \varepsilon)$ decays superexponentially as $n \rightarrow \infty$. Since $\varepsilon > 0$ was arbitrary, using (34) it follows that $d_{\text{cut}}(G_n, \kappa) \rightarrow 0$ a.s.

So far we assumed that κ was of finite type. Given an arbitrary κ , for each $\varepsilon > 0$ we can find a finite type approximation κ_ε to κ with

$$\|\kappa_\varepsilon - \kappa\|_{\text{cut}} \leq \|\kappa_\varepsilon - \kappa\|_1 \leq \varepsilon;$$

see, for example, Lemma 4.8. One can couple the random graphs $G_n = G_p(n, \kappa)$ and $G'_n = G_p(n, \kappa_\varepsilon)$ using the same vertex types x_1, \dots, x_n for each, in such a way that the symmetric difference $G_n \Delta G'_n$ has the distribution of $G_p(n, \Delta\kappa)$, where $\Delta\kappa(x, y) = |\kappa(x, y) - \kappa_\varepsilon(x, y)|$. The expected number of edges of $G_p(n, \Delta\kappa)$ is at most $n(n-1)p\|\Delta\kappa\|_1/2$ (with equality if $p\Delta\kappa \leq 1$), which is at most $n^2 p \varepsilon / 2$. It is easy to check that the actual number is tightly concentrated about the mean, so

$$d_{\text{cut}}(G_n, G'_n) \leq \|\kappa_{G_n} - \kappa_{G'_n}\|_1 = \frac{2e(G_n \Delta G'_n)}{n^2 p} \leq 2\varepsilon$$

holds with probability tending (rapidly) to 1 as $n \rightarrow \infty$. Using the finite-type case to show that $d_{\text{cut}}(G'_n, \kappa_\varepsilon) \rightarrow 0$ and the bound $d_{\text{cut}}(\kappa, \kappa_\varepsilon) \leq \varepsilon$, and recalling that $\varepsilon > 0$ was arbitrary, the result follows. \square

5 Comparison between cut and count convergence

Throughout this section, we fix a function $p = p(n)$, and consider sequences (G_n) of graphs with $|G_n| = n$. In the dense case, with $p(n) = 1$ for all n , Borgs, Chayes, Lovász, Sós and Vesztergombi [15] showed that such a sequence converges to a kernel κ in d_{cut} if and only if it converges to κ in d_{sub} ; here we wish to investigate whether this result can be extended to the sparse case. To do this, we first have to make sense of the definitions. For d_{cut} , as in the previous section, we simply associate a kernel κ_n to G_n as before, with κ_n taking the values 0 and $1/p$. Then we use the usual definition of d_{cut} for (dense) kernels to define $d_{\text{cut}}(G_n, G_m)$ and $d_{\text{cut}}(G_n, \kappa)$. In the light of Lemma 4.9, for questions of convergence the metrics d_{cut} and \widehat{d}_{cut} are equivalent; we shall use d_{cut} rather than \widehat{d}_{cut} in this section.

5.1 Admissible subgraphs and their counts

If $p = n^{-o(1)}$, then we use (17) and (18) to define d_{sub} , so convergence in d_{sub} is equivalent to convergence of $s_p(F, G_n)$ for every graph F . For smaller p , as noted in Subsection 3.5, it makes sense only to consider graphs F in a certain set \mathcal{A} of *admissible* graphs. It is not quite clear exactly which graphs should be admissible (see Subsection 3.5), so there are several variants of the definitions. To keep things simple, we shall work here with one particular choice for the set \mathcal{A} , depending on the function p . It may be that the various conjectures we shall make, if true, extend to larger sets \mathcal{A} .

Recall that we write \mathcal{F} for the set of isomorphism classes of finite (simple) graphs. Given a loopless multi-graph F and an integer $t \geq 1$, let F_t denote the graph obtained by subdividing each edge of F exactly $t - 1$ times, so $e(F_t) = te(F)$ and $|F_t| = |F| + (t - 1)e(F)$. Writing \mathcal{F}^{m} for the set of isomorphism classes of finite loopless multi-graphs, for $t \geq 2$ let

$$\mathcal{F}_t = \{F_t : F \in \mathcal{F}^{\text{m}}\},$$

and set $\mathcal{F}_1 = \mathcal{F}$ (not \mathcal{F}^{m}). Thus, for $t \geq 2$, the family \mathcal{F}_t is the set of simple graphs that may be obtained as follows: starting with a set of paths of length t , identify subsets of the endpoints of these paths in an arbitrary way, except that the two endpoints of the same path may not be identified. Note that any $F_t \in \mathcal{F}_t$ has girth at least $2t$.

Similarly, let $\mathcal{F}_{\geq t}$ be the set of simple graphs that may be obtained as above but starting with paths of length at least t . Thus $\mathcal{F}_{\geq 1} = \mathcal{F}$ and, for $t \geq 2$, $\mathcal{F}_{\geq t}$ is the set of graphs that may be obtained from some $F \in \mathcal{F}^{\text{m}}$ by subdividing each edge at least $t - 1$ times. Note that $\mathcal{F} = \mathcal{F}_{\geq 1} \supset \mathcal{F}_{\geq 2} \supset \cdots$. Let \mathcal{T} denote the set of (isomorphism classes of) finite trees.

Throughout this subsection and the next we suppose that there is some $\alpha > 0$ such that $np \geq n^\alpha$ for all large enough n . Equivalently, there is some integer $t \geq 1$ such that

$$n^{t-1}p^t \geq n^{-o(1)}. \quad (35)$$

We shall set

$$\mathcal{A} = \mathcal{T} \cup \mathcal{F}_{\geq t}$$

for the smallest such t , noting that if $p = n^{-o(1)}$ then $t = 1$, so all graphs are admissible. (An alternative that would work just as well is to let \mathcal{A} be the set of all subgraphs of graphs in $\mathcal{F}_{\geq t}$, which includes \mathcal{T} .) A key observation is that if $F \in \mathcal{F}_{\geq t}$ then (considering the internal vertices on the paths making up F) we have $|F| > e(F)(t - 1)/t$. This also holds if $F \in \mathcal{T}$, or indeed if F is a subgraph of some $F' \in \mathcal{A}$. It follows that if $F \subset F' \in \mathcal{A}$ then

$$\mathbb{E} \text{emb}(F, G(n, p)) \sim n^{|F|} p^{e(F)} = n^{|F| - e(F)(t-1)/t} (n^{t-1} p^t)^{e(F)/t} = n^{\Theta(1) - o(1)} \rightarrow \infty. \quad (36)$$

On the one hand, $\mathcal{A} = \mathcal{T} \cup \mathcal{F}_{\geq t}$ is small enough to satisfy the requirements for admissibility discussed in Subsection 3.5, including (36). (There may be requirements we have missed, in which case $\mathcal{A} = \mathcal{T} \cup \mathcal{F}_{\geq t}$ for some larger t is

likely to work.) On the other hand, as we shall now see, this set \mathcal{A} is large enough to ensure that the counts for $F \in \mathcal{A}$ determine a kernel, up to the equivalence relation \sim defined in Subsection 2.4.

Theorem 5.1. *Let κ_1 and κ_2 be two bounded kernels, and $t \geq 1$ an odd integer. Suppose that $s(F, \kappa_1) = s(F, \kappa_2)$ for every $F \in \mathcal{F}_t$. Then $\kappa_1 \sim \kappa_2$.*

Proof. Given a kernel κ , let κ^t be the kernel defined by

$$\kappa^t(x, y) = \int_{[0,1]^{t-1}} \kappa(x, x_1) \kappa(x_1, x_2) \cdots \kappa(x_{t-1}, y) dx_1 \cdots dx_{t-1}. \quad (37)$$

In other words, roughly speaking, $\kappa^t(x, y)$ counts the number of paths from x to y in κ with length t . The key observation is that if F is a graph, κ a kernel, and $t \geq 1$, then

$$s(F_t, \kappa) = s(F, \kappa^t). \quad (38)$$

Indeed, $s(F_t, \kappa)$ is defined as an integral over one variable for each vertex of F_t . We may evaluate this integral by first fixing the variables corresponding to vertices of F , then using (37) once for each edge of F to integrate over the remaining variables. What remains is exactly the integral defining $s(F, \kappa^t)$.

By assumption, $s(F, \kappa_1) = s(F, \kappa_2)$ for every $F \in \mathcal{F}_t$. Hence, from (38), we have $s(F, \kappa_1^t) = s(F, \kappa_2^t)$ for every graph F , so, by Theorem 2.8 or Theorem 2.9, $\kappa_1^t \sim \kappa_2^t$. Hence, from (12), there is a kernel κ and measure-preserving maps $\sigma_1, \sigma_2 : [0, 1] \rightarrow [0, 1]$ such that $(\kappa_i^t)^{(\sigma_i)} = \kappa$ a.e., for $i = 1, 2$. Since $(\kappa_i^{(\sigma_i)})^t = (\kappa_i^t)^{(\sigma_i)}$, we thus have $(\kappa'_1)^t = (\kappa'_2)^t$ a.e. for $\kappa'_i = \kappa_i^{(\sigma_i)}$. Since $\kappa'_i \sim \kappa_i$, and our aim is to prove that $\kappa_1 \sim \kappa_2$, it suffices to prove that $\kappa'_1 \sim \kappa'_2$. Hence, without loss of generality, we may replace κ_i by κ'_i , so we have $\kappa_1^t = \kappa_2^t$ almost everywhere. It is now a matter of simple analysis to deduce that $\kappa_1 = \kappa_2$ a.e.

Given a bounded signed kernel, i.e., a bounded function $\kappa : [0, 1]^2 \rightarrow \mathbb{R}$ satisfying $\kappa(x, y) = \kappa(y, x)$, let T_κ be the corresponding operator on $L^2([0, 1])$, defined by

$$(T_\kappa f)(x) = \int_0^1 \kappa(x, y) f(y) dy. \quad (39)$$

From the Cauchy–Schwarz inequality we have

$$\begin{aligned} \|T_\kappa f\|_2^2 &= \int_0^1 \left(\int_0^1 \kappa(x, y) f(y) dy \right)^2 dx \\ &\leq \int_0^1 \left(\int_0^1 \kappa(x, y)^2 dy \int_0^1 f(y)^2 dy \right) dx \\ &= \|f\|_2^2 \int \int \kappa(x, y)^2 dx dy = \|f\|_2^2 \|\kappa\|_2^2, \end{aligned}$$

so the operator norm of T_κ on L^2 satisfies

$$\|T_\kappa\| \leq \|\kappa\|_2 < \infty. \quad (40)$$

Now let κ be any bounded kernel, and $\varepsilon > 0$ a real number. By Lemma 4.8 there is some k such that the kernel κ_k obtained by averaging κ over 2^{-k} -by- 2^{-k} squares satisfies $\|\kappa - \kappa_k\|_2 \leq \varepsilon$. Writing $T_\kappa = T_{\kappa_k} + T_{\kappa - \kappa_k}$, the first term has finite rank, since $T_{\kappa_k} f$ is constant on intervals of length 2^{-k} . From (40), the second term has operator norm at most $\|\kappa - \kappa_k\|_2 \leq \varepsilon$. It follows that the image of the unit ball under T can be covered by a finite number of balls of radius 2ε . Since ε was arbitrary, this shows that T_κ is a compact operator.

Since κ is symmetric, we also have that T_κ self-adjoint. Consequently, T_{κ_1} is a compact self-adjoint operator on the Hilbert space $L^2(0, 1)$, so by standard results (see, for example, Bollobás [5]) there is an orthonormal basis of eigenvectors of T_{κ_1} , and all its eigenvalues are real. It is easy to see that $T_{\kappa_1^t} = (T_{\kappa_1})^t$, so $T_{\kappa_1^t}$ acts on the λ -eigenspace of T_{κ_1} by multiplication by λ^t . Since t is odd (so the map $\lambda \mapsto \lambda^t$ is injective), it follows that $T_{\kappa_1^t}$ has the same eigenspaces as T_{κ_1} . Turning this around, the action of T_{κ_1} on each eigenspace E_λ of $T_{\kappa_1^t}$ with eigenvalue λ is to multiply by $\lambda^{1/t}$. Thus, T_{κ_1} is uniquely determined by $T_{\kappa_1^t}$. In particular, since $\kappa_1^t = \kappa_2^t$ a.e., the operators T_{κ_1} and T_{κ_2} are equal, i.e., $\kappa_1 = \kappa_2$ a.e., as required. \square

Note that in Theorem 5.1 the restriction to odd t is essential, as shown by the following example.

Example 5.2. Let κ_1 and κ_2 be the two 2-by-2 ‘chessboard’ kernels defined by

$$\kappa_1(x, y) = \begin{cases} 1 & \text{if } x < 1/2, y < 1/2 \text{ or } x \geq 1/2, y \geq 1/2 \\ 0 & \text{otherwise} \end{cases}$$

and

$$\kappa_2(x, y) = \begin{cases} 1 & \text{if } x < 1/2, y \geq 1/2 \text{ or } x \geq 1/2, y < 1/2 \\ 0 & \text{otherwise} \end{cases}$$

Thus, in the dense case, κ_1 corresponds to the union of two disjoint complete graphs on $n/2$ vertices, and κ_2 to the complete $n/2$ -by- $n/2$ bipartite graph. It is easy to check that for any graph F we have $s(F, \kappa_1) = 2^{1-|F|}$, while $s(F, \kappa_2) = 2^{1-|F|}$ if F is bipartite, and $s(F, \kappa_2) = 0$ otherwise. In particular, $s(F, \kappa_1) = s(F, \kappa_2)$ for all bipartite F , and hence for all $F \in \mathcal{F}_t$, t even.

As we saw from Lemma 4.2 and Corollary 4.7, bounded density is a natural condition to impose on our sequence (G_n) when dealing with d_{cut} for sparse graphs. In the previous sections, when dealing with subgraph counts and d_{sub} , we imposed different conditions, the closest being Assumption 3.2. Let us restate this here in the appropriate form when p need not be as large as $n^{-o(1)}$.

Assumption 5.3 (exponentially bounded admissible subgraph counts).

There is a constant C such that, for each fixed $F \in \mathcal{A}$, we have $\limsup s_p(F, G_n) \leq C^{e(F)}$ as $n \rightarrow \infty$.

Note that we impose a condition only for $F \in \mathcal{A}$. When comparing d_{cut} and d_{sub} , we need to impose both Assumption 4.1 (bounded density) and Assumption 5.3. In the ‘almost dense’ case, when we take $\mathcal{A} = \mathcal{F}$, then Assumption 5.3

implies Assumption 4.1, with the same constant C . The argument is based on showing that a not-too-small dense part of G_n would contain too many $K_{t,t}$ s for some large t . Since the details are very similar to the proof of Lemma 3.5, we omit them.

Unfortunately, in general neither of Assumptions 4.1 and 5.3 implies the other. In one direction, this is easy to see: simply add a complete graph on m vertices, where $m(n)$ is chosen so that $e(K_m) \sim m^2/2 = o(n^2p)$. This does not affect Assumption 4.1, but, if m is chosen large enough, will create too many copies of any fixed connected graph F with $|F| \geq 3$. For the reverse direction, consider the following example.

Example 5.4. Fix a real number $D > 1$, and let $\kappa = \kappa_D$ be the unbounded kernel defined as follows. First partition $[0, 1]$ into intervals I_1, I_2, \dots , so that I_i has length 2^{-i} . Then set $\kappa(x, y) = i^{2/D}$ if $x, y \in I_i$, and $\kappa(x, y) = 0$ otherwise. Let F be a connected graph with average degree at most D . Then, since only terms where all vertices are in the same I_i contribute, we have

$$s(F, \kappa) = \sum_{i=1}^{\infty} 2^{-i|F|} i^{2e(F)/D} \leq \sum_{i=1}^{\infty} (2^{-i})^{|F|} \leq \sum_{i=1}^{\infty} 2^{-i} = 2.$$

Let $G_n = G_p(n, \kappa)$ be the random graph defined from κ as before. If every component of any admissible graph has average degree at most D , then it is easy to check that with probability 1 the sequence (G_n) satisfies Assumption 5.3 (with $C = 2$). On the other hand, this sequence does not satisfy Assumption 4.1, since, for every i , there will be a subgraph of G_n containing a positive fraction of the vertices with density around $i^{2/D}$.

With the choice of \mathcal{A} made here, whenever $p = p(n)$ does not satisfy $p = n^{-o(1)}$ then only trees and graphs in some $\mathcal{F}_{\geq t}$, $t \geq 2$, are admissible. All such graphs, and all their components, have average degree less than 4, so the example above shows that in this case, Assumption 5.3 does not imply Assumption 4.1.

Example 5.4 also shows that, in contrast to the almost dense case (where all graphs are admissible), in general we cannot tell from the admissible subgraph counts whether a kernel is bounded. For this reason, together with those discussed above, when comparing d_{cut} and d_{sub} we impose both Assumptions 4.1 and 5.3.

5.2 Conjectured equivalence between cut and count convergence

Our main conjecture from Section 3 was that, in the sparse case, if the subgraph counts converge, they converge to those of a kernel. In the present setting, we consider counts for admissible subgraphs. Fix $p(n)$ satisfying (35), and a set \mathcal{A} of admissible graphs. By default we take $\mathcal{A} = \mathcal{T} \cup \mathcal{F}_{\geq t}$ as in the previous subsection, although the definitions make sense for other sets \mathcal{A} . Throughout we impose Assumptions 4.1 and 5.3 for some fixed constant C . Let $X = [0, \infty)^{\mathcal{A}}$,

let $s_p : \mathcal{F} \rightarrow X$ be the map defined by

$$s_p(G_n) = (s_p(F, G_n))_{F \in \mathcal{A}} \in X$$

for any graph G_n with n vertices, let d be any metric on X inducing product topology, and define d_{sub} by mapping to X and then applying d ; as usual, we suppress the dependence on the normalizing function p . Note that d_{sub} is in general a pseudo-metric rather than a metric: there may be non-isomorphic graphs G, G' with $s_p(F, G) = s_p(F, G')$ for all $F \in \mathcal{A}$. As we only consider questions of convergence for sequences G_n with $|G_n| \rightarrow \infty$, this will not be relevant.

Let $\mathcal{L} \subset X$ denote the set of possible limit points of sequences $s_p(G_n)$, where (G_n) satisfies our assumptions.

Recall that we write \mathcal{K} for the space of kernels, that is, symmetric measurable functions $\kappa : [0, 1]^2 \rightarrow [0, C]$ quotiented by equivalence. There is a natural map from \mathcal{K} into X given by subgraph counts; we write s for this map, which does not depend on p (except through the choice of \mathcal{A}). Since \mathcal{A} always contains some set $\mathcal{F}_{\geq t}$, and hence some $\mathcal{F}_{t'}$ with t' odd, Theorem 5.1 tells us that this map is injective.

Our main conjecture is the following.

Conjecture 5.5. *With the assumptions and definitions above, we have*

$$\mathcal{L} \subset s(\mathcal{K}). \quad (41)$$

Note that if $t = 1$ then $p = n^{-o(1)}$ and we recover Conjecture 3.3; Conjecture 5.5 seems to be the natural extension of Theorem 2.1 to functions $p = p(n)$ with $p \rightarrow 0$ but $np \geq n^\alpha$ for some $\alpha > 0$.

Turning to the equivalent of Theorem 2.4, we believe that in this setting the notions of convergence given by d_{sub} and d_{cut} are equivalent. The most concrete way of saying this is as follows; again we take $\mathcal{A} = \mathcal{T} \cup \mathcal{F}_{\geq t}$ by default, although it might be that the conjecture fails for this \mathcal{A} but holds for some other \mathcal{A} .

Conjecture 5.6. *Let (G_n) be a sequence satisfying Assumptions 4.1 and 5.3, and let $\kappa \in \mathcal{K}$. Then $d_{\text{cut}}(G_n, \kappa) \rightarrow 0$ if and only if $s_p(G_n) \rightarrow s(\kappa)$.*

In this form, the conjecture implies (41) (see below). Without assuming (41), it still makes sense to compare the notions of Cauchy sequences instead.

Conjecture 5.7. *Let (G_n) be a sequence satisfying Assumptions 4.1 and 5.3. Then (G_n) is Cauchy with respect to d_{cut} if and only if (G_n) is Cauchy with respect to d_{sub} .*

As we shall shortly see, Conjectures 5.6 and 5.7 are equivalent.

Although we cannot prove the conjectures above, we can say something. Conjecture 5.6, for example, asserts two implications. Surprisingly, it is easy to show that, if (41) holds, then either of these implications (for all sequences, not just a particular sequence) implies the other! To prove this we shall first show that the random graph $G(n, \kappa)$ behaves ‘correctly’ with respect to our definition of d_{sub} ; the corresponding result for d_{cut} is Lemma 4.10.

Lemma 5.8. *Fix $C > 0$, let $\kappa : [0, 1]^2 \rightarrow [0, C]$ be a bounded kernel, and let $G_n = G_p(n, \kappa)$. Then, with probability 1, the sequence (G_n) satisfies Assumption 5.3 and we have $s_p(G_n) \rightarrow s(\kappa)$.*

Outline proof. The first statement follows from the second, since $s(F, \kappa) \leq C^{e(F)}$ holds for every F , and in particular for $F \in \mathcal{A}$.

It is well known that if F is a fixed graph, and $p' = p'(n)$ is a function of n , then the number X_F of subgraphs of $G(n, p')$ isomorphic to F is concentrated about its mean if and only if $\mathbb{E}(X_{F'}) \rightarrow \infty$ for every subgraph F' of F . (For early results of this type see Bollobás [3] and Ruciński [35]; for more recent, much stronger, results see Janson [26] and Janson, Oleszkiewicz and Ruciński [28].)

Our choice of the set \mathcal{A} ensures that this holds for every $F \in \mathcal{F}$ with $p' = Cp$ (see (36)), proving the result if κ is constant. It is straightforward to adapt this result to finite type κ . It is easy to check that for the F we consider, any $o(n^2 p)$ edges of $G_n \subset G(n, Cp)$ meet $o(n^{|F|} p^{e(F)})$ copies of F . Using this observation, one can approximate the general case by the finite type case as in the proof of Lemma 4.10. We omit the details. \square

Lemma 5.8 gives us a sequence tending in d_{sub} to any $\kappa \in \mathcal{K}$. In other words, it shows that $\mathcal{L} \supset s(\mathcal{K})$. Hence, if (41) holds,

$$\mathcal{L} = s(\mathcal{K}). \quad (42)$$

Let $\mathcal{J} \subset \mathcal{K} \times \mathcal{L}$ denote the set of pairs $(\kappa, \lambda) \in \mathcal{K} \times \mathcal{L}$ such that there is a sequence (G_n) satisfying our assumptions with $d_{\text{cut}}(G_n, \kappa) \rightarrow 0$ and $s_p(G_n) \rightarrow \lambda$. Together, Lemmas 4.10 and Lemma 5.8 tell us much more than simply that $\mathcal{L} \subset s(\mathcal{K})$: they show that the ‘diagonal’ $\mathcal{D} = \{(\kappa, s(\kappa)) : \kappa \in \mathcal{K}\}$ is contained in \mathcal{J} .

At this point, we have established three basic facts:

FACT 1: Every subsequence of (G_n) has a subsequence converging in d_{sub} to some point of \mathcal{L} . This is trivial, since Assumption 5.3 ensures that $s_p(G_n)$ lives in a compact subset of $X = [0, \infty)^{\mathcal{A}}$.

FACT 2: Every subsequence of (G_n) has a subsequence converging in d_{cut} to some kernel $\kappa \in \mathcal{K}$. This is the first part of Corollary 4.7.

FACT 3: The map s is an injection from \mathcal{K} to \mathcal{L} . As noted above, this follows from Theorem 5.1.

Facts 1 and 2 tell us that the relationship between the notions of convergence in d_{cut} and d_{sub} is described by the set \mathcal{J} . Indeed, any subsequence of (G_n) itself has a subsequence in which we have convergence in both these metrics, to some point of \mathcal{J} .

Suppose for the moment that (42) holds. There are two possibilities.

If \mathcal{J} is precisely the diagonal \mathcal{D} , then the three facts above easily imply that Conjectures 5.6 and 5.7 both hold.

If $\mathcal{J} \neq \mathcal{D}$, then there is some off diagonal point (κ_1, λ) in \mathcal{J} . Since we are assuming (42), we have $\lambda = s(\kappa_2)$ for some $\kappa_2 \in \mathcal{K}$. From the definition of \mathcal{J} there is a sequence (G_n) satisfying our assumptions, with $d_{\text{cut}}(G_n, \kappa_1) \rightarrow 0$

and $s_p(G_n) \rightarrow s(\kappa_2)$. Interleaving the sequence G_n with the sequence $G_p(n, \kappa)$, which converges to κ in both d_{cut} and d_{sub} , taking $\kappa = \kappa_1$ or κ_2 , we find a sequence which converges in one of d_{cut} or d_{sub} but not in the other. Hence, neither implication in Conjecture 5.6 or 5.7 holds, i.e., these conjectures fail as badly as possible.

In the light of the comments above, Conjecture 5.6 has the following rather vague reformulation as a question.

Question 5.9. *Given a definition of ‘suitable’ sequences (G_n) , let \mathcal{C} be the set of all graphs F with the property that, whenever κ is a bounded kernel and (G_n) is a suitable sequence with $d_{\text{cut}}(G_n, \kappa) \rightarrow 0$, then $s_p(F, G_n) \rightarrow s(F, \kappa)$. Under what reasonable definition of ‘suitable’ is the set \mathcal{C} large enough that the counts $s(F, \kappa)$, $F \in \mathcal{C}$, determine a kernel κ up to equivalence?*

The point is that, if \mathcal{C} is large enough, then the three facts above hold with $\mathcal{A} = \mathcal{C}$, and we simply use \mathcal{C} as the set of graphs whose counts we use to define d_{sub} . Then, for our ‘suitable’ sequences, d_{cut} convergence implies d_{sub} convergence to the same kernel by definition, so $(\kappa, \lambda) \in \mathcal{J}$ implies $\lambda = s(\kappa)$. Thus (41) (and hence (42)) holds, and $\mathcal{J} = \mathcal{D}$, so d_{sub} convergence also implies d_{cut} convergence. Unfortunately, there is no obvious single choice for the set of suitable sequences. One could hope that sequences with bounded density would do, but this is not the case: by adding a complete graph with many (but still $o(pn^2)$) edges to $G(n, p)$, say, it is easy to check that in this case \mathcal{C} consists only of matchings. Conjecture 5.6 is more specific than Question 5.9, since we define ‘suitable’ by assuming $s_p(F, G_n)$ bounded for F in some set \mathcal{A} , and then require $\mathcal{C} \supset \mathcal{A}$.

If Conjecture 5.5 does not hold, then Conjectures 5.6 and 5.7 cannot hold. Indeed, there is some $\lambda \in \mathcal{L}$ not corresponding to a kernel. Taking G_n converging to λ in d_{sub} , and then a subsequence that converges in d_{cut} , there is some κ with $(\kappa, \lambda) \in \mathcal{J}$. Interleaving a corresponding sequence (G_n) with $G_p(n, \kappa)$, we find a sequence that converges in d_{cut} but not in d_{sub} .

Even if Conjecture 5.5 does not hold, it is still possible that there is some relationship between cut and subgraph convergence: it may be that every sequence that is Cauchy with respect to d_{sub} , and hence converges to some $\lambda \in \mathcal{L}$, is Cauchy with respect to d_{cut} , i.e., converges to some $\kappa \in \mathcal{K}$. This happens if and only if, for every $\lambda \in \mathcal{L}$, there is a unique $\kappa \in \mathcal{K}$ such that $(\kappa, \lambda) \in \mathcal{J}$. This is not as implausible as it may sound. Indeed, suppose Conjecture 5.6 holds for some admissible set \mathcal{A}_- , but that the definitions involved make sense for a larger set \mathcal{A}_+ . It may be that (41) fails working with \mathcal{A}_+ , because we are now allowing as admissible some counts which need not converge to what we expect. However, there is a restriction map from \mathcal{L}_+ to \mathcal{L}_- forgetting about the counts outside \mathcal{A}_- . Since (42) holds for the smaller set of admissible graphs, this would show that for the larger set there is only one κ for each λ , but not vice versa.

In the next section we shall prove a form of Conjecture 5.6. Before doing so, let us briefly compare this conjecture with the corresponding result of Borgs,

Chayes, Lovász, Sós and Vesztergombi [15] for the dense case. In the dense setting, as here, Facts 1 and 2 above are easy to prove. That all limiting counts come from kernels was shown by Lovász and Szegedy [34]; this gives (42). Surprisingly, the hard part is proving Fact 3, that the counts (now meaning all counts) determine the kernel, up to equivalence as defined in Subsection 2.4. (For us this was easy, since we deduced the sparse equivalent of this statement from the dense result, Theorem 2.8.) Once one knows that the counts determine the kernel, the ‘meta-argument’ above shows that d_{cut} convergence implies d_{sub} convergence if and only if the reverse implication holds. Since the forward implication is very easy (see Corollary 2.3), the result of [15] that the two metrics are equivalent follows. This gives a proof of this result in which the only non-straightforward step is showing that the counts $s(F, \kappa)$ determine the kernel κ up to the appropriate notion of equivalence. One might expect this uniqueness result to be easy, but this seems to be far from the case. Recently, Borgs, Chayes and Lovász [12] gave a direct proof of this result (which, as noted in Section 2, actually follows from the results of [15]); their proof is far from simple.

5.3 Partial results: embedding lemmas

Our aim in this section is to prove a positive result, that under certain circumstances, if $d_{\text{cut}}(G_n, \kappa) \rightarrow 0$, then $s_p(F, G_n) \rightarrow s_p(F, \kappa)$ for certain graphs F . In the case where κ is of finite type, this is simply a counting lemma: in this case, $G_n \rightarrow \kappa$ says that G_n can be partitioned into (ε, p) -regular pairs with densities given by κ . In the uniform case, Chung and Graham [17] proved such counting lemmas for certain graphs under certain assumptions. The general case turns out to be rather different, but we shall still use several of their ideas.

We start with the simplest case, where F is a path. First we need some definitions. As usual, in the proof it will be easier to consider homomorphisms from F to G_n (i.e., walks in G_n) rather than embeddings. As we shall see later, this makes no difference.

For G_n a graph and X_0, \dots, X_ℓ subsets of $V(G_n)$, let $G_n(X_0, X_1, \dots, X_\ell)$ denote the number of $(\ell + 1)$ -tuples (v_i) with $v_i \in X_i$ and $v_i v_{i+1} \in E(G)$ for $0 \leq i \leq \ell - 1$. Identifying a subset of $V(G_n)$ with a subset of $[0, 1]$ as before, for a kernel κ let

$$\kappa(X_0, X_1, \dots, X_\ell) = \int_{X_0 \times \dots \times X_\ell} \kappa(x_0, x_1) \cdots \kappa(x_{\ell-1}, x_\ell) dx_0 \cdots dx_\ell.$$

Lemma 5.10. *Let $C > 0$ be constant, let $p(n)$ be any function of n with $np \rightarrow \infty$, and let (G_n) be a sequence of graphs with $t_p(T, G_n)$ bounded for each tree T . For every $\varepsilon > 0$ and $\ell \geq 1$ there is a $\delta = \delta_\ell(\varepsilon) > 0$ such that, whenever $\kappa : [0, 1]^2 \rightarrow [0, C]$ is a kernel with $\|\kappa_{G_n} - \kappa\|_{\text{cut}} \leq \delta$, then*

$$|G_n(X_0, X_1, \dots, X_\ell) - n^{\ell+1} p^\ell \kappa(X_0, X_1, \dots, X_\ell)| \leq \varepsilon n^{\ell+1} p^\ell$$

for any sets $X_0, X_1, \dots, X_\ell \subset V(G_n)$.

Roughly speaking, the lemma says that if $G_n \rightarrow \kappa$ and $t_p(T, G_n)$ is bounded for each T , then $t_p(P_\ell, \kappa) \rightarrow s(P_\ell, \kappa)$. The stronger assertion makes it simpler to prove the result by induction.

Proof. Renormalizing, we may assume without loss of generality that $C = 1$. Let us do so from now on.

The fact that δ is not allowed to depend on κ allows us to assume without loss of generality that κ is piecewise constant on squares of side $1/n$, i.e., that κ may be interpreted as a (dense) weighted graph with vertex set $V(G_n)$. Indeed, the Frieze–Kannan form of Szemerédi’s Lemma shows that there is an integer k such that, given any κ , there is a κ' that is constant on squares of side $1/k$ with $d_{\text{cut}}(\kappa, \kappa') \leq \delta$. Tweaking κ' slightly if k does not divide n , we obtain a kernel κ'' of the required form. Replacing δ by 2δ as appropriate, the result for κ follows from the result for κ'' . [Note that we implicitly assumed that n is large here, meaning larger than some n_0 depending on ε and ℓ . We could simply assume this in the statement of the lemma, but it can be achieved by subdividing vertices. In fact, we could work with a kernel instead of a graph throughout the proof.]

Let

$$\Delta(X_0, \dots, X_\ell) = \frac{G_n(X_0, X_1, \dots, X_\ell)}{n^{\ell+1}p^\ell} - \kappa(X_0, X_1, \dots, X_\ell),$$

so our aim is to show that $|\Delta(X_0, \dots, X_\ell)| \leq \varepsilon$ for all choices of the sets X_i . We shall show much more: let $M = \max_T \sup_n t_p(T, G_n)$, where the maximum is over trees with at most $2\ell + 1$ vertices, noting that $M < \infty$. We shall show that if $d_{\text{cut}}(G_n, \kappa) \leq \delta$, then, for any $1 \leq t \leq \ell$ and any $X_0, \dots, X_t \subset V(G_n)$ we have

$$|\Delta(X_0, X_1, \dots, X_t)| \leq \varepsilon_t, \quad (43)$$

where $\varepsilon_1 = \delta$, and

$$\varepsilon_t = 7\sqrt{\varepsilon_{t-1}} + \sqrt{M}\varepsilon_{t-1}^{1/4}$$

for $t \geq 2$. Since ε_ℓ tends to zero as $\delta \rightarrow 0$, taking δ small enough we have $\delta = \varepsilon_1 \leq \varepsilon_2 \leq \dots \leq \varepsilon_\ell \leq \varepsilon$, so to complete the proof of the lemma it suffices to prove (43) for this choice of δ .

We shall prove (43) by induction on t . For $t = 1$, the result is immediate from the definition of the cut norm: indeed, $\Delta(X_0, X_1)$ is one of the quantities appearing in the supremum defining this norm. Suppose now that $2 \leq t \leq \ell$, and that (43) holds with t replaced by $t - 1$.

For $v \in V(G)$ and $X_1, \dots, X_r \subset V(G)$, set

$$\kappa(v, X_1, \dots, X_r) = \int_{X_1 \times \dots \times X_r} \kappa(x, x_1) \kappa(x_1, x_2) \cdots \kappa(x_{r-1}, x_r) dx_1 \cdots dx_r,$$

where x is any point of the interval of length $1/n$ corresponding to the vertex v , and let

$$\Delta(v, X_1, \dots, X_r) = \frac{G_n(\{v\}, X_1, \dots, X_r)}{n^r p^r} - \kappa(v, X_1, \dots, X_r). \quad (44)$$

Note that

$$\Delta(X, X_1, \dots, X_r) = \frac{1}{n} \sum_{v \in X} \Delta(v, X_1, \dots, X_r). \quad (45)$$

Fix $X_0, \dots, X_t \subset V(G_n)$, and set $\eta = \sqrt{\varepsilon_{t-1}}$. Let B_1 be the set of $v \in X_1$ with $\Delta(v, X_2, \dots, X_t) > \eta$. Then, from (45), $\Delta(B_1, X_2, \dots, X_t) \geq \eta|B_1|/n$. But by the induction hypothesis, $\Delta(B_1, X_2, \dots, X_t) \leq \varepsilon_{t-1} = \eta^2$. Hence, $|B_1| \leq \eta n$. Arguing similarly, and using $\varepsilon_{t-1} \geq \varepsilon_1$, we see that the set B of vertices $v \in X_1$ for which either $|\Delta(v, X_2, \dots, X_t)| \geq \eta$ or $|\Delta(v, X_0)| \geq \eta$ holds has size at most $4\eta n$.

If $v \in X_1 \setminus B$, then we have roughly the right number of walks through v , i.e.,

$$G_n(X_0, \{v\}, X_2, \dots, X_t) = G_n(\{v\}, X_0)G_n(\{v\}, X_2, \dots, X_t)$$

is close to $np\kappa(v, X_0)n^{t-1}p^{t-1}\kappa(v, X_2, \dots, X_t)$. More precisely, using the fact that κ is pointwise bounded by $C = 1$ to bound the κ terms in the last expression by 1, for $v \in X_1 \setminus B$ we have

$$|\Delta(X_0, v, X_2, \dots, X_t)| \leq 3\eta, \quad (46)$$

where the left hand side is defined by analogy with (44).

It remains to consider $v \in B$. For $i = 1, 2$, let

$$\sigma_i = \sum_{v \in B} G_n(X_0, \{v\}, X_2, \dots, X_t)^i,$$

noting that $\sigma_1 \leq \sqrt{|B|\sigma_2}$ by the Cauchy-Schwarz inequality. Let T be the tree with $2t$ edges formed by identifying the second vertices of two paths of length t . Then σ_2 counts a subset of the homomorphisms from T into G_n , so

$$\sigma_2 \leq \text{hom}(T, G_n) = n^{2t+1}p^{2t}t_p(T, G_n) \leq Mn^{2t+1}p^{2t}.$$

Since $|B| \leq 4\eta n$ it follows that

$$\sigma_1 \leq \sqrt{|B|\sigma_2} \leq 2\sqrt{M\eta}n^{t+1}p^t.$$

Since κ is bounded by 1, we have $\kappa(X_0, B, X_2, \dots, X_t) \leq \mu(B) \leq 4\eta$, so

$$|\Delta(X_0, B, X_2, \dots, X_t)| \leq 2\sqrt{M\eta} + 4\eta.$$

Together with the bound (46) for $v \in X_1 \setminus B$ and (the equivalent of) (45), this implies that

$$|\Delta(X_0, X_1, \dots, X_t)| \leq 7\eta + 2\sqrt{M\eta} = \varepsilon_t,$$

as required. This completes the proof of (43) by induction, and thus the proof of the lemma. \square

Note that the argument above works just as well for an arbitrary fixed tree rather than a path: we pick some leaf v to play the role of x_0 ; the unique neighbour of v then plays the role of x_1 . This gives us a counting lemma for trees.

Corollary 5.11. *Let (G_n) be a sequence of graphs with $t_p(T, G_n)$ bounded for every tree T , and suppose that $d_{\text{cut}}(G_n, \kappa) \rightarrow 0$, where κ is a bounded kernel. Then for each tree T we have $t_p(T, G_n) \rightarrow s(T, \kappa)$ as $n \rightarrow \infty$. \square*

Chung and Graham [17] proved a version of this result (for paths rather than trees) with κ constant, under the assumption that the maximum degree of G_n is at most Cpn . This maximum degree assumption of course gives $t_p(T, G_n) \leq C^{e(T)}$, so it is stronger than the bounded tree counts assumption of Lemma 5.10. In some sense, the maximum degree condition is much stronger, but it turns out that our global assumption is just as good for questions involving subgraph counts. The reason that Lemma 5.10 is more complicated than the corresponding simple result in [17] is that κ is not uniform, not our weaker assumption.

We stated earlier that, in the sparse case, the parameter $s_p(F, \kappa)$ should be preferred to $t_p(F, \kappa)$, even though t_p tends to be easier to work with. Nevertheless, in the case of trees, these parameters are equivalent, as shown by the following observation.

Lemma 5.12. *Let $p(n)$ be any function of n with $np \rightarrow \infty$, and let (G_n) be a sequence with $s_p(T, G_n)$ bounded for every tree T . Then, for each tree T , we have $t_p(T, G_n) \sim s_p(T, G_n)$. In particular, $t_p(T, G_n)$ is bounded.*

Proof. Fix a tree T with k vertices. It suffices to show that the number N_T of non-injective homomorphisms from T to G_n satisfies $N_T = o(n^k p^{k-1})$ as $n \rightarrow \infty$. Now the image of any non-injective homomorphism ϕ from T to G_n is a connected subgraph H of G_n with ℓ vertices, where $1 \leq \ell \leq k-1$. Any such subgraph contains a tree T' with ℓ vertices, so for each ℓ there are (crudely) at most $\sum_{|T'|=\ell} \text{emb}(T', G_n)$ possibilities for vertex set of H , where the sum is over all trees T' with ℓ vertices. Since there are at most k^ℓ homomorphisms ϕ with image a given set of ℓ vertices, we thus have

$$N_T \leq \sum_{\ell=1}^{k-1} k^\ell \sum_{|T'|=\ell} \text{emb}(T', G_n).$$

Since $\text{emb}(T', G_n) = n_{(|T'|)} p^{e(T')} s_p(T', G_n)$, the final term is $O(n^\ell p^{\ell-1})$ by assumption. It follows that $N_T = O(n^{k-1} p^{k-2}) = o(n^k p^{k-1})$, as claimed. \square

Lemma 5.12 allows us to restate Corollary 5.11 in terms of the parameter s .

Theorem 5.13. *Let (G_n) be a sequence of graphs with $s_p(T, G_n)$ bounded for every tree T , and suppose that $d_{\text{cut}}(G_n, \kappa) \rightarrow 0$, where κ is a bounded kernel. Then for each tree T we have $s_p(T, G_n) \rightarrow s(T, \kappa)$ as $n \rightarrow \infty$. \square*

Theorem 5.13 may be regarded as an embedding lemma for trees. Our next aim is to prove a much more general result. Chung and Graham showed that, in the uniform case, if the number of paths of length $\ell - 1$ between any two vertices is at most a constant times what it should be, then almost all pairs of

vertices are joined by almost the right number of paths of length ℓ , and hence G_n contains asymptotically the expected number of copies of any $F \in \mathcal{F}_\ell$. This result is much harder than the paths result, even in the uniform case. Although we shall use the key idea of Chung and Graham, the proof does not carry over in a simple way. In the following result, we work with t_p rather than s_p for simplicity; we return to this later.

Theorem 5.14. *Let $C > 0$ and $\ell \geq 3$ be fixed, and let $p = p(n)$ be any function of n . Let (G_n) be a sequence of graphs with $\sup_n t_p(F, G_n) < \infty$ for each $F \in \mathcal{T} \cup \mathcal{F}_\ell \cup \{C_{2\ell-2}\}$, and suppose that $d_{\text{cut}}(G_n, \kappa) \rightarrow 0$ for some kernel $\kappa : [0, 1]^2 \rightarrow [0, C]$. Then $t_p(F, G_n) \rightarrow s(F, \kappa)$ for each $F \in \mathcal{F}_\ell$.*

Proof. Note that by Lemma 4.2, the sequence (G_n) has density bounded by C , i.e., it satisfies Assumption 4.1. Renormalizing, we shall assume without loss of generality that $C = 1$.

Fix $\varepsilon > 0$, and a graph $F_\ell \in \mathcal{F}_\ell$. Let $\eta > 0$ be a small constant to be chosen below (depending on ε , ℓ and F_ℓ). By Lemma 4.5 there is some K such that for n large enough, which we assume from now on, G_n has an (η, p) -regular partition $\Pi = (P_1, \dots, P_k)$ for some $k = k(n) \leq K$. Passing to a subsequence of (G_n) , we may assume that k is constant. As usual, we shall ignore rounding to integers, assuming that each P_i contains exactly n/k vertices.

Passing to a subsequence (again), we may assume that for all i and j the sequence $d_p(P_i, P_j)$ converges to some $\kappa'(P_i, P_j) \in [0, 1]$. Relabelling if necessary so that P_i consists of vertices v with $in/k < v \leq (i+1)n/k$, and identifying vertices with corresponding subsets of $[0, 1]$ as usual, we may view κ' as a kernel on $[0, 1]^2$.

If n is large enough, which we assume, then each $d_p(P_i, P_j)$ is within ε of $\kappa(P_i, P_j)$. It follows that $d_{\text{cut}}(G_n/\Pi, \kappa') \leq \|G_n/\Pi - \kappa'\|_1 \leq \varepsilon$. Under our bounded density assumption 4.1, strong regularity implies weak regularity (for suitably transformed parameters), so choosing η small enough we have $d_{\text{cut}}(G_n, G_n/\Pi) \leq \varepsilon$. Hence, choosing n large enough, $d_{\text{cut}}(\kappa, \kappa') \leq d_{\text{cut}}(\kappa, G_n) + d_{\text{cut}}(G_n, G_n/\Pi) + d_{\text{cut}}(G_n/\Pi, \kappa') \leq 3\varepsilon$. Hence, by Lemma 2.2, for any fixed F we have

$$|s(F, \kappa) - s(F, \kappa')| = O(\varepsilon),$$

so it suffices to show that $t_p(F_\ell, G_n)$ is close to $s(F_\ell, \kappa')$ rather than to $s(F_\ell, \kappa)$. To avoid clutter in the notation, from now on we write κ for the finite type kernel κ' defined above; the original κ plays no further role in the proof. Recall that κ (formerly known as κ') is bounded by 1. For $u \in P_i$ and $v \in P_j$ we shall abuse notation by writing $\kappa(u, v) = \kappa(P_i, P_j)$ for the value of κ at any point of $[0, 1]^2$ corresponding to (u, v) . Recall that $|d_p(P_i, P_j) - \kappa(P_i, P_j)| \leq \varepsilon$ for all i, j .

For $v, w \in V(G_n)$ and $t \geq 1$, let $w_t(v, w)$ denote the number of walks of length t in G_n starting at v and ending at w ; we suppress the dependence on G_n in the notation. Let $\kappa^t(v, w)$ denote the normalized ‘expected’ number of such walks, if G_n behaved like the random graph $G_p(n, \kappa)$. Let $U \subset V^2$ be the set of pairs (v, w) such that $w_\ell(v, w) \leq (\kappa^\ell(v, w) - \varepsilon)n^{\ell-1}p^\ell$. We call the pairs $(v, w) \in U$ *underconnected*, since they are joined by ‘too few’ walks of length ℓ .

We shall show that

$$|U| = |\{(v, w) : w_\ell(v, w) \leq (\kappa^\ell(v, w) - \varepsilon)n^{\ell-1}p^\ell\}| \leq \varepsilon n^2 \quad (47)$$

if η is chosen suitably, and then n is taken large enough. Before doing so, let us note that this implies the result.

By Lemma 5.10, if we choose η small enough, then the total number of walks of length ℓ in G_n is within $\varepsilon n^{\ell+1}p^\ell$ of the expected number in $G_p(n, \kappa)$, namely $\|\kappa^\ell\|_1 n^{\ell+1}p^\ell$. If (47) holds, then if we count only a maximum of $\kappa^\ell(v, w)n^{\ell-1}p^\ell$ walks for each pair (v, w) of endpoints, we still count at least $(1 - \varepsilon)(\|\kappa^\ell\|_1 - \varepsilon)n^{\ell+1}p^\ell$ walks, so there are at most $3\varepsilon n^{\ell+1}p^\ell$ walks uncounted, using $\|\kappa^\ell\|_1 \leq 1$. Writing W for the set of *overconnected* pairs $(v, w) \in V^2$ with $w_\ell(v, w) \geq (\kappa^\ell(v, w) + \sqrt{\varepsilon})n^{\ell-1}p^\ell$, it follows that

$$|W| \leq 3\sqrt{\varepsilon}n^2. \quad (48)$$

In other words, almost all pairs of vertices are joined by almost the right number of walks.

Recall that we fixed a graph $F_\ell \in \mathcal{F}_\ell$. Let F_ℓ be obtained by subdividing the edges of a loopless multi-graph F with vertex set u_1, \dots, u_r , so

$$\text{hom}(F_\ell, G_n) = \sum_{v_1, \dots, v_r \in V(G_n)} \prod_{u_i u_j \in E(F)} w_\ell(v_i, v_j), \quad (49)$$

where the factors in the product corresponding to multiple edges of F are of course repeated. Given $u_i u_j \in E(F)$, let $2F_\ell/E_2$ be the graph formed from two copies of F_ℓ by identifying the vertices corresponding to u_i and identifying the vertices corresponding to u_j . Since $2F_\ell/E_2 \in \mathcal{F}_\ell$, we have $t_p(2F_\ell/E_2, G_n)$ bounded. It follows by the Cauchy–Schwarz inequality that the number of homomorphisms from F_ℓ into G_n mapping u_i and u_j to a pair in $U \cup W$ is small, in fact of order $\varepsilon^{1/4}n^{|F_\ell|}p^{e(F_\ell)}$; the argument is as in the proof of Lemma 5.10.

Since the comment above applies to any edge $u_i u_j$ of F , the contribution to the sum in (49) from terms in which one or more pairs (v_i, v_j) fall in $U \cup W$ is small. But in the remaining terms, $w_\ell(v_i, v_j)$ is well approximated by $\kappa^\ell(v_i, v_j)n^{\ell-1}p^\ell$, and it follows that $t_p(F_\ell, G_n)$ is close to $s(F_\ell, \kappa)$: the difference is bounded by some function of $|F_\ell|$ and ε . In short, we have shown that to prove the theorem, it suffices to prove (47), i.e., that there are few underconnected pairs.

From now on, we forget the original graph F_ℓ , and aim to prove (47), recalling that κ is a fixed finite-type kernel and that G_n/Π is (pointwise) within ε of κ , where $\Pi = (P_1, \dots, P_k)$ is our (η, p) -regular partition of G_n . It will be convenient to assume that ε is fairly small. In particular, we shall assume that $\varepsilon \leq 1/40$.

Recall that all but at most ηk^2 pairs in our partition $(P_i)_1^k$ are (η, p) -regular. Since *all* pairs have density at most $1 + \varepsilon \leq 2$, the irregular pairs contain at most $2\eta n^2 p$ edges. By assumption $t_p(T, G_n)$ is bounded for each tree T , and in particular for the trees formed from two paths by identifying an edge from each, so using Cauchy–Schwarz again a small set of edges meets only a small fraction

of the walks of length ℓ in G_n . In particular, the number of walks of length ℓ containing one or more edges from irregular pairs is $O(\sqrt{\eta}n^{\ell+1}p^\ell)$. Taking η small enough, we may assume that this quantity is less than $\varepsilon^2 n^{\ell+1}p^\ell/10$, say. It follows that in proving (47), we may delete all edges in irregular pairs, i.e., we may assume that every pair is regular: if (47) holds for the resulting graph G'_n and kernel κ' with $\varepsilon/2$ in place of ε , then (47) holds for our original graph G_n and kernel κ .

The lower bound in the proof of Lemma 5.10 used only closeness of the graph and kernel in the cut norm, not the bounds on various tree counts. This argument can thus be applied locally to sequences of parts of our partition. Abusing notation, let us write $P_0, P_1, \dots, P_{\ell-1}$ for an arbitrary sequence of ℓ parts of our partition, with repetition allowed. For any subsets $X_i \subset P_i$, we find that there are at least

$$p^{\ell-1} \prod_{i=0}^{\ell-1} |X_i| \prod_{i=0}^{\ell-2} \kappa(P_i, P_{i+1}) - \gamma \frac{n^\ell}{k^\ell} p^{\ell-1}$$

walks $v_0 v_1 \dots v_{\ell-1}$ with $v_i \in X_i$, where $\gamma = \gamma(\eta, \ell)$ tends to 0 as $\eta \rightarrow 0$. We choose η small enough that $\gamma \leq \varepsilon^{12}$. Taking $X_i = P_i$ for $i > 0$, and summing over all choices for the intermediate parts, a consequence of this is that if P_0 and $P_{\ell-1}$ are any two parts, and X_0 is any subset of P_0 , then there are at least

$$(\kappa^{\ell-1}(P_0, P_{\ell-1})|X_0|/|P_0| - \gamma)n^\ell p^{\ell-1}/k^2 \quad (50)$$

walks of length $\ell - 1$ from X_0 to $P_{\ell-1}$.

Let us call a walk of length $\ell - 1$ in G_n *bad* if there are at least $Mn^{\ell-2}p^{\ell-1}$ walks in G_n with the same endpoints, where M is a constant to be chosen in a moment, depending on ε but not on η ; otherwise, the walk is *good*. Each bad walk may be extended to at least $Mn^{\ell-2}p^{\ell-1}$ homomorphic images of $C_{2\ell-2}$. By assumption, $t_p(C_{2\ell-2}, G_n)$ is bounded, so it follows that there are $O(n^\ell p^{\ell-1}/M)$ bad walks. In particular, choosing the constant M large enough, we may assume that there are at most $\varepsilon^9 n^\ell p^{\ell-1}/3$ bad walks.

Suppose for a contradiction that (47) does not hold, i.e., the set U of underconnected pairs of vertices has size at least εn^2 . Our first aim is to select a pair (P, P') of parts of our partition such that there are many underconnected pairs (u, v) in $P \times P'$, but not too many bad walks start in P . Since $|U| \geq \varepsilon n^2$ by assumption, there are at least $\varepsilon k/2$ parts P with

$$|U \cap (P \times V)| \geq \varepsilon n^2/(2k). \quad (51)$$

On the other hand, there are at most $\varepsilon k/3$ parts P with the property that more than $\varepsilon^8 n^\ell p^{\ell-1}/k$ bad walks start in P (otherwise there would be too many bad walks). Hence there exists a part P for which (51) holds, with at most $\varepsilon^8 n^\ell p^{\ell-1}/k$ bad walks starting in P . Fix such a P . From (51) and averaging, there is a part P' such that

$$|U \cap (P \times P')| \geq \varepsilon n^2/(2k^2) = \varepsilon |P||P'|/2. \quad (52)$$

From now on, fix such a P' .

Let us say that a pair (u, P'') with $u \in P$ and P'' a part of our partition is *deficient* if there are fewer than $(\kappa^{\ell-1}(P, P'') - \sqrt{\gamma})n^{\ell-1}p^{\ell-1}/k$ walks of length $\ell - 1$ from u to P'' , where γ is as in (50). For a given P'' , at most $\sqrt{\gamma}n/k$ vertices $u \in P$ form a deficient pair with P'' : otherwise, the set X_0 of such vertices would have more than $\gamma n^{\ell}p^{\ell-1}/k^2$ fewer walks to P'' than it should have, contradicting (50). Hence, there are at most $\sqrt{\gamma}n$ deficient pairs. Let $D \subset P$ be the set of vertices u in more than $\gamma^{1/4}k$ deficient pairs. Then $|D| \leq \sqrt{\gamma}n/(\gamma^{1/4}k) = \gamma^{1/4}|P|$.

Let us say that a pair (u, P'') with $u \in P$ and P'' a part of our partition is *compromised* if there are more than $\varepsilon^3 n^{\ell-1}p^{\ell-1}/k$ bad walks from u to P'' . Since at most $\varepsilon^8 n^{\ell}p^{\ell-1}/k$ bad walks start in P , there are at most $\varepsilon^5 n$ compromised pairs. Let C be the set of $u \in P$ in more than $\varepsilon^3 k$ compromised pairs; then $|C| \leq \varepsilon^2 n/k = \varepsilon^2 |P|$.

Let $S \subset P$ be the set of vertices u for which there are at least $\varepsilon|P'|/4$ vertices $v \in P'$ with $(u, v) \in U$. By (52) we have

$$\varepsilon|P||P'|/2 \leq |U \cap (P \times P')| \leq |S||P'| + \varepsilon|P||P'|/4,$$

so $|S| \geq \varepsilon|P|/4 > (\gamma^{1/4} + \varepsilon^2)|P|$. Thus $|S| > |D| + |C|$, and there is some u in $S \setminus (D \cup C)$. Fix such a u for the rest of the proof, and let U_u denote the set of $v \in P'$ for which (u, v) is underconnected.

At this point we have chosen a vertex $u \in P$, a part P' , and a set $U_u \subset P'$ with the following properties:

- (i) for each $v \in U_u$, there are at most $(\kappa^{\ell}(u, v) - \varepsilon)n^{\ell-1}p^{\ell} = (\kappa^{\ell}(P, P') - \varepsilon)n^{\ell-1}p^{\ell}$ walks of length ℓ from u to v .
- (ii) $|U_u| \geq \varepsilon|P'|/4$,
- (iii) there are at most $\gamma^{1/4}k \leq \varepsilon^3 k$ deficient pairs (u, P'') ,
- (iv) there are at most $\varepsilon^3 k$ compromised pairs (u, P'') .

From (i) and (ii) above, there are at least $m = \varepsilon^2 n^{\ell}p^{\ell}/(4k)$ ‘missing walks’ from u to U_u : the number of walks of length ℓ from u to U_u falls short of the expected number in $G_p(n, \kappa)$ by at least m . Let P'' be any part of our partition. By a u - U_u walk via P'' we mean a walk of length ℓ from u to U_u whose second last vertex lies in P'' ; the expected number of such walks is $N_{P''} = \kappa^{\ell-1}(P, P'')\kappa(P'', P')|U_u|n^{\ell-1}p^{\ell}/k$. Note that $\sum_{P''} N_{P''}$ is simply the expected number of walks from u to U_u . Let $m_{P''}$ be the number of ‘missing walks via P'' ’, i.e., the difference between $N_{P''}$ and the number of u - U_u walks via P'' , or zero if there are at least $N_{P''}$ such walks. The total number of missing walks is at most the sum of the numbers $m_{P''}$, so

$$\sum_{P''} m_{P''} \geq m \geq \varepsilon^2 n^{\ell}p^{\ell}/(4k).$$

Let us say that P'' is *useful* if $m_{P''} \geq \varepsilon^2 n^{\ell}p^{\ell}/(8k^2)$, so the contribution to the sum above from non-useful parts P'' is at most half the right hand side. Recalling that we have normalized so that κ is bounded by 1, and that $\varepsilon < 1/40$,

for each P'' we have $m_{P''} \leq N_{P''} \leq n^\ell p^\ell / k^2$; it follows that there are at least $\varepsilon^2 k / 8 \geq 5\varepsilon^3 k$ useful parts P'' .

Using (iii) and (iv) above, it follows that there is a part P'' which is useful, but neither deficient nor compromised. Fix such a part P'' .

Recall that a walk of length $\ell - 1$ from u to $w \in P''$ is *good* if it is not bad, i.e., if

$$w_{\ell-1}(u, w) \leq N = Mn^{\ell-2}p^{\ell-1}. \quad (53)$$

Since $\gamma^{1/4} \leq \varepsilon^3$, and P'' is neither deficient nor compromised, there are at least

$$(\kappa^{\ell-1}(P, P'') - 2\varepsilon^3)n^{\ell-1}p^{\ell-1}/k$$

good walks from u to P'' . On the other hand, there are many missing walks via P'' . With this setup, we are finally ready to apply the key idea of Chung and Graham [17], which is to partition the set P'' into subsets according to the approximate number of walks from u to the relevant vertex, and then use regularity to show that there are about the right number of walks from U_u to each such subset. In fact, there is a slick way of doing this.




Figure 1: The set P'' is subdivided into sets A_i , with $w_{\ell-1}(u, v) = i$ for each $v \in A_i$. Each edge from A_i to U_u contributes i walks from u to U_u via P'' .

For $i \geq 0$, let A_i be the set of vertices $v \in P''$ with $w_{\ell-1}(u, v) = i$; see Figure 1. Also, let $A_i^+ = \bigcup_{j \geq i} A_j$. Then,

$$w_{\ell-1}(u, P'') = \sum_{i \geq 0} i|A_i| = \sum_{i \geq 1} |A_i^+|.$$

More importantly, $\sum_{i=1}^N |A_i^+|$ is at least the number of *good* walks from u to P'' , so

$$\sum_{i=1}^N |A_i^+| \geq (\kappa^{\ell-1}(P, P'') - 2\varepsilon^3)n^{\ell-1}p^{\ell-1}/k. \quad (54)$$

Since (P'', P') is (η, p) -regular with (normalized) density $\kappa(P'', P') \leq 1$, if $A \subset P''$ and $B \subset P'$ then $e(A, B) \geq p\kappa(P'', P')|A||B| - \eta p(n/k)^2$ (this is trivially true if one of A or B has size less than $\eta n/k$). Since each edge from U_u to A_i forms the final edge of exactly i walks from u to U_u , the number of walks from u to U_u via P'' is given by

$$\begin{aligned} \sum_{i \geq 1} i e(A_i, U_u) &\geq \sum_{i=1}^N e(A_i^+, U_u) \\ &\geq \sum_{i=1}^N p\kappa(P'', P')|A_i^+||U_u| - \eta p(n/k)^2 \\ &\geq (\kappa^{\ell-1}(P, P'') - 2\varepsilon^3)\kappa(P'', P')|U_u|n^{\ell-1}p^\ell/k - \eta Np(n/k)^2, \end{aligned}$$

where we used (54) in the last step. The main term is simply the expected number of walks from u to U_u via P'' , so the conclusion is that there are at most

$$2\varepsilon^3\kappa(P'', P')|U_u|n^{\ell-1}p^\ell/k + \eta Np(n/k)^2 \quad (55)$$

missing walks from u to U_u via P'' . The two terms above may be bounded above by $2\varepsilon^3n^\ell p^\ell/k^2$ and, recalling (53), $\eta Mn^\ell p^\ell/k^2$, respectively. Choosing $\eta \leq \varepsilon^3/M$ we thus have at most $3\varepsilon^3n^\ell p^\ell/k^2$ missing walks via P'' , i.e., $m_{P''} \leq 3\varepsilon^3n^\ell p^\ell/k^2$, which contradicts the fact that P'' is useful. This contradiction completes the proof. \square

Note that the argument above does not extend to $\ell = 2$, and not only because C_2 makes no sense. The problem is that we cannot define N as in (53) (this quantity is now $o(1)$), but must take $N = 1$ instead, and then the second term in (55) is too large.

The proof of Theorem 5.14 actually gives rather more with almost no extra work.

Theorem 5.15. *Let $C > 0$ and $\ell \geq 3$ be fixed, and let $p = p(n)$ be any function of n . Let (G_n) be a sequence of graphs with $\sup_n t_p(F, G_n) < \infty$ for each $F \in \mathcal{T} \cup \mathcal{F}_{\geq \ell} \cup \{C_{2\ell-2}\}$, and suppose that $d_{\text{cut}}(G_n, \kappa) \rightarrow 0$ for some bounded kernel κ . Then $t_p(F, G_n) \rightarrow s(F, \kappa)$ for each $F \in \mathcal{T} \cup \mathcal{F}_{\geq \ell}$.*

Proof. The conclusion for $F \in \mathcal{T}$ follows from Corollary 5.11.

Fix $F \in \mathcal{F}_{\geq \ell}$ and $\varepsilon > 0$, and let L be the length of the longest induced path in F . Noting that for $t > \ell$ we have $C_{2t-2} \in \mathcal{F}_{\geq \ell}$, the hypotheses of Theorem 5.14 are satisfied with ℓ replaced by any t in the range $\ell \leq t \leq L$. The proof of that result thus shows that if η is chosen small enough, then when we take an (η, p) -regular partition of G_n with associated kernel κ' , almost all pairs (v, w) of vertices are joined by almost the ‘right’ number of walks of each length t , $\ell \leq t \leq L$. More precisely, writing κ for κ' as in the proof of Theorem 5.14, and writing U_t for the set of pairs (v, w) with $w_t(v, w) \leq (\kappa^t(v, w) - \varepsilon)n^{t-1}p^t$ and W_t for the set of pairs with $w_t(v, w) \geq (\kappa^t(v, w) + \sqrt{\varepsilon})n^{t-1}p^t$, the proof of Theorem 5.14 shows that $|U_t| \leq \varepsilon n^2$ for $\ell \leq t \leq L$, and (hence) that $|W_t| \leq$

$3\sqrt{\varepsilon}n^2$ for each t in this range. Using the analogue of (49) in which each term $w_\ell(\cdot, \cdot)$ is replaced by an appropriate term $w_t(\cdot, \cdot)$, as before we can use the Cauchy–Schwarz inequality to show that the contribution to $t_p(F, G_n)$ from terms with some pair (v_i, v_j) in the small set $\bigcup_t U_t \cup W_t$ is small (of order $\varepsilon^{1/4}$), and it follows as before that if η is small enough, then $|t_p(F, G_n) - s(F, \kappa)|$ is bounded by some function of F and ε , giving the result. \square

Let us note for later reference that, in one way, the assumptions of Theorems 5.14 and 5.15 are weaker than they may first appear. Let F be a loopless multigraph with vertex set u_1, u_2, \dots, u_k , and let $F_\ell \in \mathcal{F}_\ell$ be obtained by subdividing each edge of F exactly $\ell - 1$ times. Then (49) may be rewritten as

$$\text{hom}(F_\ell, G_n) = n^k \mathbb{E} \left(\prod_{u_i u_j \in E(F)} w_\ell(v_i, v_j) \right),$$

where the expectation is over the uniform choice of $(v_1, v_2, \dots, v_k) \in V(G_n)^k$. Applying Hölder’s inequality, $\mathbb{E}(\prod_{i=1}^r X_i) \leq (\prod \mathbb{E}(|X_i|^r))^{1/r}$, with $r = e(F)$, it follows that

$$\begin{aligned} \text{hom}(F_\ell, G_n)^r &\leq n^{kr} \prod_{u_i u_j \in E(F)} \mathbb{E}(w_\ell(v_i, v_j)^r) = n^{kr} \mathbb{E}(w_\ell(v_1, v_2)^r)^r \\ &= n^{kr-2r} \text{hom}(H_{r,\ell}, G_n)^r, \end{aligned} \quad (56)$$

where $H_{r,\ell} \in \mathcal{F}_\ell$ is the ‘theta graph’ consisting of r internally vertex disjoint paths of length ℓ joining the same pair of vertices. The normalizing factors work out correctly, so we have

$$t_p(F_\ell, G_n) \leq t_p(H_{r,\ell}, G_n). \quad (57)$$

Hence, the condition that $t_p(F, G_n)$ remain bounded for every $F \in \mathcal{F}_\ell$ is equivalent to the condition that $t_p(F, G_n)$ is bounded for $F = H_{r,\ell}$, $r = 1, 2, \dots$

Arguing similarly, for any $F \in \mathcal{F}_{\geq \ell}$ we may bound $t_p(F, G_n)$ in terms of the quantities $t_p(H_{r,\ell'}, G_n)$, where ℓ' ranges over the lengths of the paths making up F . Hence, to show that $t_p(F, G_n)$ is bounded for all $F \in \mathcal{F}_{\geq \ell}$, it suffices to prove the same condition for the graphs $H_{r,\ell'}$, $r \geq 1$, $\ell' \geq \ell$. Note that these latter conditions are simply moment conditions on the numbers of walks of various lengths joining a random pair of vertices of G_n .

In the case where the limiting kernel κ is of finite type, Theorem 5.15 may be seen as a form of counting lemma. In this case, it is easy to strengthen the result to count homomorphisms from F into G_n with each vertex mapped to a specified part of the partition of G_n corresponding to the finite type kernel κ , obtaining a result similar in form to Lemma 5.10. Such a (strengthened) finite type case of Theorem 5.14 or Theorem 5.15 is very much easier to prove than the general case: there is no need to apply Szemerédi’s Lemma, and the proof of the result of Chung and Graham [17] for the uniform case goes through without

much modification. One might hope that, using Szemerédi's Lemma, the full generality of Theorem 5.15 would follow easily from the finite type case, but this is not true. The problem is that our assumptions are inescapably global: we assume, for instance, that the number of copies of $C_{2\ell-2}$ in G_n is bounded by a multiple of the expected number of copies. When we take an (ε, p) -regular partition, this gives no useful information about the number of copies of $C_{2\ell-2}$ in each regular pair: we have a bound that is of the form $Mk^{2\ell-2}$ times the expected number of copies, where k is the number of parts. To apply the finite type case, we would need a bound independent of k . For this reason there seems to be no easy way around the work in the proof of Theorem 5.14.

Theorem 5.15 may be seen as some progress towards a proof of some form of Conjecture 5.6. More precisely, it is almost an answer to Question 5.9: the only problem is that for Theorem 5.15 we work with t_p rather than s_p . We shall return to this in detail in a moment. However, even ignoring this, Theorem 5.15 is a little disappointing in some ways. Let $\mathcal{A} = \mathcal{T} \cup \mathcal{F}_{\geq \ell}$. Assuming boundedness of $t_p(F, G_n)$ for $F \in \mathcal{A} \cup \{C_{2\ell-2}\}$, we obtain convergence of the counts $t_p(F, G_n)$ for $F \in \mathcal{A}$. The extra assumption for $F = C_{2\ell-2}$ is somehow annoying. This is perhaps clearest if we consider the range where p is fairly large, say $n^{-o(1)}$. In this case $s_p \sim t_p$, and it makes sense to assume boundedness of all counts $s_p(F, G_n)$. However, since C_2 does not make sense, the smallest value of ℓ for which we can apply Theorem 5.15 is $\ell = 3$, and we obtain convergence of the counts $s_p(F, G_n)$ for $F \in \mathcal{F}_{\geq 3}$. In comparison, Theorem 3.20 shows that with the counts s_p bounded, and $s_p(C_4, G_n) \rightarrow 1$, which should roughly correspond to convergence to the uniform kernel $\kappa = 1$, we obtain $s_p(F, G_n) \rightarrow s(F, \kappa) = 1$ for all $F \in \mathcal{F}_{\geq 2}$, rather than just for $F \in \mathcal{F}_{\geq 3}$.

In fact, Theorem 3.20 gives much more: it gives convergence for all F with girth at least 4. Chung and Graham [17] asked whether an analogous result holds for sparse graphs under the appropriate assumptions (what they call ' ℓ -quasi randomness', which corresponds roughly to the assumptions of Theorem 5.14 with κ constant), with girth at least 4 replaced by girth at least 2ℓ . In our language, they asked whether (when $\kappa = 1$) the conclusion of Theorem 5.14 can be extended to all F with girth at least 2ℓ . Unfortunately, the answer is no for a trivial reason, namely that there are graphs F with arbitrarily large girth and arbitrarily large average degree. Taking $p = n^{-\alpha}$ for some $0 < \alpha < 1$, and d large enough, for any graph F with average degree d the expected number of copies of F in $G_n = G(n, p)$ is $o(1)$, so the normalizing constant in the definition of $t_p(F, G_n)$ is $o(1)$. Since $\text{hom}(F, G_n)$ is an integer, we cannot have $t_p(F, G_n) \rightarrow 1$ in this case.

5.4 Embeddings or homomorphisms?

In this subsection we return to the use of t_p rather than s_p in Theorems 5.14 and 5.15. Although this simplifies the proof, it is unsatisfactory for a reason we shall now explain. We start by discussing the analogous problem with the corresponding result of Chung and Graham [17], their Theorem 8. We shall use

the following fact, proved by Blakley and Roy [2] in a slightly more general form in the context of symmetric matrices.

Theorem 5.16. *Let G be a graph with n vertices and average degree d . Then G contains at least nd^ℓ walks of length ℓ .* \square

Recall that we write $w_t(u, v)$ for the number of walks of length t from u to v . Chung and Graham [17] impose the condition that $w_{\ell-1}(u, v) < c_0 p^{\ell-1} n^{\ell-2}$ holds for *every* pair of vertices u, v , where c_0 is a constant: they call this *condition* $U(\ell)$. In other words, the number of walks from u to v is at most a constant times what it should be. Normalizing so that G_n contains exactly $pn^2/2$ edges, Chung and Graham note that $U(\ell)$ can only hold if $p = \Omega(n^{-1+1/(\ell-1)})$: otherwise, the expected number of walks of length $\ell-1$ from a random u to a random v is much less than 1, so $w_{\ell-1}(u, v)$ must sometimes be much larger than its expectation.

In fact, $U(\ell)$ cannot hold unless p is quite a bit larger, but for the ‘wrong’ reason: taking ℓ odd for simplicity, let $\ell = 2k + 1$. Considering walks of length $\ell-1$ formed by tracing a walk of length k forwards and then backwards, we see that if G_n has $pn^2/2$ edges, then

$$\sum_v w_{\ell-1}(v, v) \geq \text{hom}(P_k, G_n) \geq n(np)^k, \quad (58)$$

where the second inequality is Theorem 5.16. Thus there is some v with $w_{\ell-1}(v, v) \geq (np)^k$, and it follows that $U(\ell)$ can only hold if $p = \Omega(n^{-1+2/(\ell-1)})$, so Theorem 8 of [17] can only be applied for p in this range. Note that this is an essential problem: this result counts homomorphisms (Chung and Graham use the notation $\#\{H \subset G\}$ for $\text{hom}(H, G)$), and the bound on $w_{\ell-1}(u, v)$ is definitely used with $u = v$. Indeed, as we shall see, the conclusion fails if $p = o(n^{-1+2/(\ell-1)})$.

Turning to Theorem 5.14, the condition that $t_p(C_{2\ell-2}, G_n)$ remain bounded corresponds roughly to the condition $U(\ell)$: indeed, the former says exactly that

$$\sum_{u,v} w_{\ell-1}(u, v)^2 = O(n^{2\ell-2} p^{2\ell-2}), \quad (59)$$

which follows immediately from $U(\ell)$. It turns out that the problem described above does not arise with (59) – in this second moment (rather than uniform) condition, the few pairs with $u = v$ matter less. Indeed, it is easy to check that in $G(n, p)$, for example, (59) holds as long as $p = \Omega(n^{-1+1/(\ell-1)})$. [The expected number of homomorphisms from $C_{2\ell-2}$ whose image is a tree with k edges is $O(n(np)^k) = O(n(np)^{\ell-1})$, and the expected number whose image is a graph with k vertices containing a cycle is $O(n^k p^k) = O((np)^{2\ell-2})$.] However, the same problem arises in a different place.

As before, let $H_{k,\ell} \in \mathcal{F}_\ell$ be the ‘theta graph’ formed by k paths of length ℓ joining the same pair (s, t) of vertices, with the paths internally vertex disjoint. Suppose that ℓ is even. Writing $w_t(v) = w_t(v, V(G_n))$ for the number of walks of length t in G_n starting at v , normalizing still so that $e(G_n) = pn^2/2$, and

considering homomorphisms from $H_{k,\ell}$ to G_n mapping s and t to a common vertex v , we have

$$\text{hom}(H_{k,\ell}, G_n) \geq \sum_v w_{\ell/2}(v)^k \geq n \left(\frac{1}{n} \sum_v w_{\ell/2}(v) \right)^k \geq n(np)^{k\ell/2},$$

where the second inequality is from convexity and the last from Theorem 5.16. Since $|H_{k,\ell}| = 2 + k(\ell - 1)$ and $e(H_{k,\ell}) = k\ell$, it follows that $t_p(H_{k,\ell}, G_n) \geq n^{k-1}(np)^{-k\ell/2}$. Suppose that $p \leq n^{-1+2/\ell-\varepsilon}$ for some $\varepsilon > 0$. Then taking k large enough we see that $t_p(H_{k,\ell}, G_n) \rightarrow \infty$, so neither the assumptions nor the conclusion of Theorem 5.14 can hold. When ε is small, this value of p is much larger than that above which the number of subgraphs of $G(n, p)$ isomorphic to $H_{k,\ell}$ is well behaved.

The calculations above illustrate the problem with working with t_p : we count certain trees as copies of $H_{k,\ell}$, for example, and the number of these trees exceeds the number of embeddings of $H_{k,\ell}$ in a wide range of densities in which Theorem 5.14 might otherwise apply. For this reason, if we could replace t_p by s_p throughout the statement of the theorem, we would obtain a much stronger and more satisfactory result: not only would it count embeddings, which is what we are really interested in, but it would apply to a much larger family of graphs, for example, to random graphs with much lower densities. Unfortunately, the proof breaks down in various places if we simply replace t_p by s_p . However, the next result is a major step in this direction.

Given vertices v, w of a graph G_n , suppressing the dependence on G_n , let us write $p_\ell(v, w)$ for the number of *paths* of length ℓ from v to w , so $p_\ell(v, w) \leq w_\ell(v, w)$.

Theorem 5.17. *Let $C > 0$ and $\ell \geq 3$ be fixed, and let $p = p(n)$ be any function of n . Let (G_n) be a sequence of graphs satisfying the following three conditions:*

$$\sup_n s_p(F, G_n) < \infty \text{ for each } F \in \mathcal{T}, \quad (60)$$

$$\sum_u \sum_{v \neq u} p_{\ell-1}(u, v)^2 = O(n^{2\ell-2} p^{2\ell-2}), \quad (61)$$

and

$$\sum_u \sum_{v \neq u} p_\ell(u, v)^k = O(n^{2+k(\ell-1)} p^{k\ell}), \quad (62)$$

for each fixed $k \geq 1$. Suppose also that $d_{\text{cut}}(G_n, \kappa) \rightarrow 0$ for some kernel $\kappa : [0, 1]^2 \rightarrow [0, C]$. Then $s_p(F, G_n) \rightarrow s(F, \kappa)$ for each $F \in \mathcal{F}_\ell$.

Before turning to the proof of this result, let us make some remarks on the conditions above. Firstly, in (60) it makes no difference whether we write s_p or t_p , by Lemma 5.12.

Condition (62) is almost the same as the condition $s_p(H_{k,\ell}, G_n) = O(1)$. Indeed, $\text{emb}(H_{k,\ell}, G_n)$ is simply the sum over distinct u and v of the number of

k -tuples of *internally vertex disjoint* paths from u to v , so (62), which bounds the same sum without the restriction to disjoint paths, is formally stronger than $s_p(H_{k,\ell}, G_n) = O(1)$. Since there are (typically) many paths from u to v in the range of p for which (61) may hold, it seems very likely that, assuming the other conditions of Theorem 5.17, $s_p(H_{k,\ell}, G_n) = O(1)$ implies (62), so (62) could be replaced by this more pleasant condition. However, we do not have a proof of this.

Similarly, condition (61) is closely related to $s_p(C_{2\ell-2}, G_n) = O(1)$, and could perhaps be replaced by this weaker condition. This is less clear, however, as Theorem 5.17 can be applied for p small enough that the typical number of paths of length $\ell - 1$ between a given pair of vertices is $O(1)$.

Instead of (61) we can always impose the stronger condition $t_p(C_{2\ell-2}, G_n) = O(1)$; these conditions are probably equivalent in the present setting. The corresponding statement for (62) and the stronger assumption $t_p(H_{k,\ell}, G_n) = O(1)$ is not true; see the discussion of the behaviour of $t_p(H_{k,\ell}, G_n)$ in the paragraphs preceding Theorem 5.17.

Finally, let us note that (62) gives us control over $s_p(F_\ell, G_n)$ for all $F_\ell \in \mathcal{F}_\ell$, not just for $F_\ell = H_{k,\ell}$. Let F_ℓ be obtained by subdividing a graph F with vertex set u_1, u_2, \dots, u_k . Then

$$\text{emb}(F_\ell, G_n) \leq \sum_{v_1, v_2, \dots, v_k} \prod_{u_i u_j \in E(F)} p_\ell(v_i, v_j),$$

where the sum is over all $n_{(k)}$ k -tuples of distinct vertices of G_n . Applying Hölder's inequality as in the proof (56) of (57), but in a probability space with $n_{(k)}$ elements rather than n^k , we find that

$$\text{emb}(F_\ell, G_n) \leq n_{(k)} \mathbb{E}(p_\ell(v_1, v_2)^{e(F)}),$$

where the expectation is over the choice of a random pair (v_1, v_2) of *distinct* vertices of G_n . Condition (62) bounds the final expectation; as usual the normalizing factors work out, and we see that if (62) holds for every k then $s_p(F_\ell, G_n) = O(1)$ for every $F_\ell \in \mathcal{F}_\ell$.

Outline proof of Theorem 5.17. Since the proof is a relatively simple modification of that of Theorem 5.14, we shall give only an outline, concentrating on the differences.

The first change we make is that we work with paths rather than walks, replacing the quantities $w_t(u, v)$, $t = \ell - 1, \ell$, appearing in the proof of Theorem 5.14 with the corresponding quantities $p_t(u, v)$. By Lemma 5.12, all but a vanishing fraction of the walks in G_n of a given length are paths, so (47), for example, implies the same statement with $w_\ell(v, w)$ replaced by $p_\ell(v, w)$. Of course, (47) was proved using the assumption $t_p(C_{2\ell-2}, G_n) = O(1)$, whereas we now have the weaker assumption (61). However, following through the proof it is easy to see that if we count paths instead of walks, then (61) suffices. (The key point is that (61) suffices to bound the number of *bad paths*, i.e., paths between endpoints u, v with $p_{\ell-1}(u, v) > Mn^{\ell-2}p^{\ell-1}$.)

Let us fix (a small) $\varepsilon > 0$ and a graph $F_\ell \in \mathcal{F}_\ell$. We also fix an integer N to be chosen later, depending only on ε and F_ℓ . Finally, let η be a small positive constant depending on ε , F_ℓ and N . For reasons that will become clear later, we first partition $V(G_n)$ into N almost equal parts Q_1, \dots, Q_N . Then we take an (η, p) -regular partition (P_i) with each P_i contained in some Q_j . For the moment we ignore the partition (Q_i) .

As before, passing to a subsequence we assume that the densities $d_p(P_i, P_j)$ converge to a finite-type kernel κ . Let $S \subset V \times V$ be the set of pairs of vertices joined by the ‘wrong’ number of paths of length ℓ :

$$S = \{(v, w) : v \neq w, |p_\ell(v, w) - \kappa^\ell(v, w)n^{\ell-1}p^\ell| \geq \varepsilon n^{\ell-1}p^\ell\}.$$

If η is chosen small enough then the proofs of (47) and (48) carry through counting paths instead of walks, and (replacing ε by $\varepsilon^2/10$), the equivalents of (47) and (48) imply that

$$|S| \leq \varepsilon n(n-1). \quad (63)$$

We proceed from here to our bound on $s_p(F_\ell, G_n)$ in two steps. First we count something that is not quite an embedding of F_ℓ .

Let F_ℓ be obtained from the loopless multigraph F by subdividing each edge $\ell - 1$ times, and let u_1, \dots, u_k be the vertices of F , which we also regard as vertices of F_ℓ . By a *semiembedding* of F_ℓ into G_n we mean a homomorphism from F_ℓ into G_n that maps the vertices u_1, \dots, u_k to distinct vertices of G_n , and each of the $e(F)$ u_i - u_j paths of length ℓ that make up the graph F_ℓ into a *path* in G_n . Clearly, every embedding is a semiembedding; the only additional condition on an embedding is that the paths in G_n are internally vertex disjoint.

Let $\text{emb}^+(F_\ell, G_n) \geq \text{emb}(F_\ell, G_n)$ denote the number of semiembeddings of F_ℓ into G_n . Then, from the definition of a semiembedding, we have

$$\text{emb}^+(F_\ell, G_n) = \sum_{v_1, \dots, v_k} \prod_{u_i u_j \in E(F)} p_\ell(v_i, v_j), \quad (64)$$

where the sum is over all $n_{(k)}$ sequences (v_1, \dots, v_k) of distinct vertices of G_n and, as usual, any multiple edges in F give rise to multiple factors in the product.

As before we, we can rewrite the formula above as an expectation over a random choice of (v_1, \dots, v_k) . Normalizing correctly for a change, let X_{ij} be the random variable $p_\ell(v_i, v_j)/(n^{\ell-1}p^\ell)$, so

$$s_p^+(F_\ell, G_n) = \frac{\text{emb}^+(F_\ell, G_n)}{n_{(|F_\ell|)}p^{e(F_\ell)}} \sim \frac{\text{emb}^+(F_\ell, G_n)}{n_{(|F|)}n^{|F_\ell|-|F|}p^{e(F_\ell)}} = \mathbb{E} \left(\prod_{u_i u_j \in E(F)} X_{ij} \right).$$

Equation (63) says, roughly speaking, that each X_{ij} is with high probability close to ‘what it should be’, which is a random variable depending on κ , the kernel corresponding to the partition (P_1, \dots, P_k) of G_n . We should like to deduce that the expectation of the product is close to what it should be.

Let Z be the set of k -tuples (v_1, \dots, v_k) with the v_i distinct such that $(v_i, v_j) \in S$ for some $1 \leq i < j \leq k$. Regarding Z as an event in our probability space,

$$\mathbb{P}(Z) \leq \binom{k}{2} \mathbb{P}((v_1, v_2) \in S) \leq \varepsilon \binom{k}{2},$$

from (63). Hölder's inequality thus gives

$$\mathbb{E} \left(1_Z \prod_{u_i u_j \in E(F)} X_{ij} \right) \leq \left(\mathbb{E} \left(1_Z^{e(F)+1} \right) \prod_{u_i u_j \in E(F)} \mathbb{E} \left(X_{ij}^{e(F)+1} \right) \right)^{1/(e(F)+1)},$$

where 1_Z is the indicator function of the event Z . Now, for each i and j , we have

$$\mathbb{E}(X_{ij}^{e(F)+1}) = \frac{1}{n(n-1)} \sum_u \sum_{v \neq u} \left(\frac{p_\ell(u, v)}{n^{\ell-1} p^\ell} \right)^{e(F)+1},$$

which is $O(1)$ by our assumption (62). Also, $\mathbb{E}(1_Z^{e(F)+1}) = \mathbb{E}(1_Z) = \mathbb{P}(Z) \leq \varepsilon$. Hence,

$$\mathbb{E} \left(1_Z \prod_{u_i u_j \in E(F)} X_{ij} \right) = O(\varepsilon^{1/(e(F)+1)}). \quad (65)$$

In other words, the contribution to (64) from semiembeddings mapping some edge of F into a pair $(u, v) \in S$ is negligible. By definition of S , the contribution from all other semiembeddings is ‘what it should be’, and it follows that

$$|s_p^+(F_\ell, G_n) - s(F_\ell, \kappa)| \leq O(\varepsilon^{1/(e(F)+1)}) + O(\varepsilon).$$

Since $\varepsilon > 0$ was arbitrary, we thus have $s_p^+(F_\ell, G_n) \sim s(F_\ell, \kappa)$.

In the end, of course, it is $s_p(F_\ell, G_n)$ that we wish to bound, not $s_p^+(F_\ell, G_n)$. Since $s_p(F_\ell, G_n) \leq s_p^+(F_\ell, G_n)$ it remains to show that most semiembeddings are in fact embeddings, i.e., that the paths in G_n making up a typical semiembedding are internally vertex disjoint. For paths corresponding to vertex disjoint edges of F , this is quite easy, using the fact that $s_p(T, G_n)$ is bounded for each tree, which tells us that almost all pairs of paths of length ℓ are vertex disjoint. For paths corresponding to edges of F sharing a vertex, there is a similar argument. We shall not spell these arguments out as there is a third case that cannot be handled in this way, namely paths corresponding to duplicate edges in F . We must allow these, since we include, for example, $C_{2\ell}$ in F_ℓ . It is in handling these paths that our ‘crude’ partition (Q_i) comes in.

Let us classify paths $w_0 w_1, \dots, w_\ell$ in G_n into $N^{\ell+1}$ types, according to which part Q_i each w_i lies in. We say that a pair (u, v) of distinct vertices of G_n is *good* if, for all $N^{\ell-1}$ possible types of u - v path, the number of u - v paths of this type is ‘close’ to what it should be, i.e., within $\varepsilon |Q_1|^{\ell-1} p^\ell \sim \varepsilon n^{\ell-1} p^\ell / N^{\ell-1}$ of what it should be. As usual, ‘what it should be’ means the expected number in $G_p(n, \kappa)$, which depends not only on which parts P_i the vertices u and v lie in,

but also on the type of path being considered. Let S' be the set of pairs (u, v) , $u \neq v$, that are *bad*, i.e., not good.

Since N is fixed before η is chosen, it is not hard to see that the argument giving (63) (applied with $\varepsilon/N^{\ell-1}$ in place of ε) also shows that $|S'| \leq \varepsilon n(n-1)$; we omit the details. In other words, almost all pairs of vertices are joined by about the right number of paths of any given type. As before, we break down the set of embeddings of F_ℓ into G_n according to which vertices v_1, \dots, v_k of G_n the ‘branch vertices’ u_1, \dots, u_k are mapped to. Defining Z' analogously to Z , but using S' instead of S , the argument giving (65) shows that we may assume that $(v_1, \dots, v_k) \notin Z'$, i.e., that no pair (v_i, v_j) is in S' . Counting embeddings with v_1, \dots, v_k fixed, it remains to choose $e(F)$ paths joining the appropriate pairs v_i, v_j . Let us choose these paths one by one. Since the total number of paths joining v_i to v_j is about what it should be, all we must show is that few (say at most $\varepsilon n^{\ell-1} p^\ell$) paths from v_i to v_j meet one of our at most $e(F) - 1$ earlier paths. But this is now easy: we must avoid a set X of at most $(e(F) - 1)(\ell - 1) = O(1)$ vertices, the internal vertices of the previously chosen paths. In fact, we shall do much more, avoiding any part Q_a that meets X ! This rules out at most $(\ell - 1)|X|N^{\ell-2}$ of the $N^{\ell-1}$ types of v_i-v_j paths. Choosing N large enough (larger than $1/\varepsilon$), this is only a fraction $O(\varepsilon)$ of all possible types. Since $(v_i, v_j) \notin S'$, we have almost the right number of paths of each remaining type, and hence almost the right number of paths in total. This completes our outline proof of Theorem 5.17. \square

Of course, there is a variant of Theorem 5.17 which is to Theorem 5.17 as Theorem 5.15 is to Theorem 5.14; we shall not state this separately.

Let us close this section by giving one simple example of a setting in which the conditions of Theorem 5.17 are satisfied. Fix $\ell \geq 3$, and suppose that our sequence (G_n) has the following two properties. Firstly, the maximal degree $\Delta(G_n)$ is not too large:

$$\Delta(G_n) \leq Mpn, \quad (66)$$

for some constant M . Secondly,

$$p_{\ell-1}(u, v) \leq Mn^{\ell-2}p^{\ell-1} \quad (67)$$

for all $u \neq v \in V(G_n)$. Condition (66) is called DEG in Chung and Graham [17]; condition (67) is related to their condition $U(\ell)$, but, as noted in the paragraph containing (58), is much weaker. In particular, it is easy to check that if $p = n^{-\alpha}$ with $0 < \alpha < 1$ constant, and κ is any bounded kernel, then the random graphs $G_p(n, \kappa)$ satisfy (66) and (67) with probability 1, as long as $\alpha < 1 - 1/(\ell - 1)$. If (66) and (67) hold then $p_\ell(v, w) \leq M^2 n^{\ell-1} p^\ell$ for all v and w , while $s_p(T, G_n) \leq M^{e(T)}$ for any tree T , so the conditions of Theorem 5.17 are satisfied. Similarly, $p_t(v, w) \leq M^{t-\ell+2} n^{t-1} p^t$ holds for all $t \geq \ell$, so the variant of Theorem 5.17 corresponding to Theorem 5.15 applies.

It follows that conditions (66) and (67) provide an answer to Question 5.9. Indeed, Theorem 5.17 tells us that, under these conditions, if κ is a bounded kernel, then $d_{\text{cut}}(G_n, \kappa) \rightarrow 0$ implies $s_p(F, G_n) \rightarrow s(F, \kappa)$ for all $F \in F_\ell$; its

variant gives us $s_p(F, G_n) \rightarrow s(F, \kappa)$ for all $F \in F_{\geq \ell}$. By Theorem 5.1, the counts $s(F, \kappa)$, $F \in F_{\geq \ell}$, do determine the kernel (up to equivalence), so conditions (66) and (67) are ‘suitable’ in the sense of Question 5.9. As noted after Question 5.9, this implies the following result.

Theorem 5.18. *Fix $\ell \geq 3$, let $p = p(n)$ be any function, and let (G_n) be a sequence of graphs satisfying (66), (67) and the bounded density assumption 4.1. Then, for any bounded kernel κ , we have $d_{\text{cut}}(G_n, \kappa) \rightarrow 0$ if and only if $d_{\text{sub}}(G_n, \kappa) \rightarrow 0$, where d_{sub} is defined using $\mathcal{A} = \mathcal{T} \cup \mathcal{F}_{\geq \ell}$ for the set of admissible graphs. \square*

In this section we discussed how to extend the subgraph (count) metric to sparse graphs, noting that there are various possibilities (depending on the choice of the set \mathcal{A} of admissible graphs), and conjectured that one particular extension is equivalent to the cut metric. In the next section we turn to a different metric, that extends much more easily to sparse graphs.

6 The partition metric

As noted in Section 2, for dense graphs there are many natural metrics that turn out to be equivalent, in the sense of generating the same topology. So far we have focussed on the cut and subgraph (or count) metrics; we now turn to the *partition metric*, introduced by Borgs, Chayes, Lovász, Sós and Vesztergombi [16]. In the dense case, it turns out to be relatively easy to show that the partition and cut metrics are equivalent; in this brief section we show that, under mild assumptions, this equivalence holds also in the sparse setting, as long as $np \rightarrow \infty$.

On the one hand, this result (Theorem 6.2, below) shows that for graphs with $\omega(n)$ edges, no new questions arise by considering the partition metric. On the other hand, it reinforces the conclusion that the cut metric remains extremely natural for sparse graphs, and gives a way of considering the cut metric from a very different point of view. There is another, very important, motivation for introducing partition metrics for sparse graphs: when we come to extremely sparse graphs, with $\Theta(n)$ edges, the cut metric turns out to make very little sense, while the partition metric (which is no longer equivalent) remains natural. This is a major topic in its own right and will be discussed in a companion paper [11].

6.1 Partition matrices and the partition metric

Turning to the formal definitions, as in the rest of the paper, let $p = p(n)$ be a normalizing function and G_n a graph with n vertices. Let $k \geq 2$ be fixed. For $n \geq k$ and $\Pi = (P_1, \dots, P_k)$ a partition of $V(G_n)$ into k non-empty parts, let $M_\Pi(G_n) = (d_p(P_i, P_j))_{1 \leq i, j \leq k}$ be the matrix encoding the normalized densities of edges between the parts of Π (see (27)). Since $M_\Pi(G_n)$ is symmetric, we may think of this matrix as an element of $\mathbb{R}^{k(k+1)/2}$. Set

$$\mathcal{M}_k(G_n) = \{M_\Pi(G_n)\} \subset \mathbb{R}^{k(k+1)/2},$$

where Π runs over all balanced partitions of $V(G_n)$ into k parts, i.e., all partitions (P_1, \dots, P_k) with $|P_i| - |P_j| \leq 1$.

As usual, we assume that G_n has $O(pn^2)$ edges. For definiteness, let us assume that $e(G_n) \leq Cpn^2/2$. Since each part of a balanced partition has size at least $n/(2k)$, the entries of any $M_\Pi(G_n) \in \mathcal{M}_k(G_n)$ are bounded by $C_k = (2k)^2 C$, say. Thus, $\mathcal{M}_k(G_n)$ is a subset of the compact space $B_k = [0, C_k]^{k(k+1)/2}$.

Let $\mathcal{C}_0(B_k)$ denote the set of non-empty compact subsets of B_k , and let d_H be the Hausdorff metric on $\mathcal{C}_0(B_k)$, defined with respect to the ℓ_∞ distance, say. Thus

$$d_H(X, Y) = \inf\{\varepsilon > 0 : X^{(\varepsilon)} \supset Y, Y^{(\varepsilon)} \supset X\},$$

where $X^{(\varepsilon)}$ denotes the ε -neighbourhood of X in the ℓ_∞ metric. Since (B_k, ℓ_∞) is compact, by standard results (see, for example, Dugundji [20, p. 253]), the space $(\mathcal{C}_0(B_k), d_H)$ is compact. To ensure that the metric we are about to define is a genuine metric, it is convenient to work with $\mathcal{C}(B_k) = \mathcal{C}_0(B_k) \cup \{\emptyset\}$, setting $d_H(\emptyset, X) = C_k$, say, for any $X \in \mathcal{C}(B_k)$, so the empty set is an isolated point in $(\mathcal{C}(B_k), d_H)$.

Let $\mathcal{C} = \prod_{k \geq 2} \mathcal{C}(B_k)$, and let $\mathcal{M} : \mathcal{F} \mapsto \mathcal{C}$ be the map defined by

$$\mathcal{M}(G_n) = (\mathcal{M}_k(G_n))_{k=2}^\infty$$

for every graph G_n on n vertices, noting that $\mathcal{M}_k(G_n)$ is empty if $k > n$. Then we may define the *partition metric* d_{part} by

$$d_{\text{part}}(G, G') = d(\mathcal{M}(G), \mathcal{M}(G')),$$

where d is any metric on \mathcal{C} giving rise to the product topology. Considering the partition of an n vertex graph into n parts shows that d_{part} is a metric on the set \mathcal{F} of isomorphism classes of finite graphs. Recalling that each space $(\mathcal{C}(B_k), d_H)$ is compact, the key property of the partition metric is that (G_n) is Cauchy with respect to d_{part} if and only if there are compact sets $Y_k \subset B_k$ such that $d_H(\mathcal{M}_k(G_n), Y_k) \rightarrow 0$ for each k . In particular, convergence in d_{part} is equivalent to convergence of the set of partition matrices for each fixed k . Thus we may always think of k as fixed and n as much larger than k .

In the dense case, a metric equivalent to d_{part} has been introduced independently by Borgs, Chayes, Lovász, Sós and Vesztegombi [16]; the only difference is that in [16], all partitions into k parts are considered, rather than just balanced partitions. Of course, one then needs to take care to ensure that the densities between small parts are counted with an appropriate weight when computing the distance between density matrices \mathcal{M}_k . Whether one takes all partitions or just balanced partitions is a matter of taste: it is very easy to see that convergence in either of the resulting metrics implies convergence in the other.

We may extend the map $\mathcal{M} : \mathcal{F} \rightarrow \mathcal{C}$, and hence d_{part} , to bounded kernels in a natural way: instead of partitioning the vertex set into k almost equal parts, we partition $[0, 1]$ into k exactly equal parts, and consider the closure of the

set of ‘density matrices’ that may be obtained from κ using such partitions; we omit the details. Note that, as shown by Borgs, Chayes, Lovász, Sós and Vesztergombi [16, Example 4.4], the set of density matrices is not in general closed.

As for the cut metric, it is easy to check that it makes little difference whether we define d_{part} for graphs directly, or by going via kernels. (The corresponding dense result appears in [16]: the sparse case here is slightly more complicated due to the possibility of ‘high-degree’ vertices.)

Lemma 6.1. *Let $p = p(n)$ satisfy $p \geq 1/n$, and let (G_n) be a sequence of graphs with $e(G_n) = O(pn^2)$ and $\Delta(G_n) = o(pn^2)$. Then $d_{\text{part}}(G_n, \kappa_{G_n}) \rightarrow 0$ as $n \rightarrow \infty$.*

Proof. By definition, we must show that $d_H(\mathcal{M}_k(G_n), \mathcal{M}_k(\kappa_{G_n})) \rightarrow 0$ for each $k \geq 1$. Fix k . Since $e(G_n) = O(pn^2)$, there is a constant D such that at most $n/(2k)$ vertices of G_n have degree more than Dpn . Let L denote the set of ‘low-degree’ vertices, with degree at most Dpn , so $|L| \geq n - n/(2k)$.

We must show that for any density matrix in $\mathcal{M}_k(G_n)$ there is a nearby matrix in $\mathcal{M}_k(\kappa_{G_n})$, and vice versa. The forward implication is trivial: a balanced partition Π of $V(G_n)$ corresponds to a partition of $[0, 1]$ into sets whose sizes differ by $O(1/n) = o(1)$. Adjusting these parts slightly, making changes only in subintervals of $[0, 1]$ corresponding to low degree vertices, the entries of the corresponding density matrix change by $o(1)$.

For the reverse implication, let Π be a partition of $[0, 1]$ into k parts P_1, \dots, P_k , and let $M \in \mathcal{M}_k(\kappa_{G_n})$ be the corresponding density matrix, with entries m_{ij} . For $v \in V(G_n) = [n]$ and $1 \leq i \leq k$, let $p_{v,i}$ be the fraction of the subinterval of $[0, 1]$ corresponding to the vertex v that lies in P_i , noting that $\sum_i p_{v,i} = 1$ for each v , and $\sum_v p_{v,i} = n/k$ for each i . Form a random partition $\Pi' = (P'_1, \dots, P'_k)$ as follows: put each vertex v into a random part P'_{i_v} with $\mathbb{P}(i_v = i) = p_{v,i}$, with the choices independent for different vertices v .

It is immediate that $\mathbb{E}(|P'_i|) = n/k$ and $\text{Var}(|P'_i|) \leq n/k$. It follows that for some constant C we have

$$\forall i : ||P'_i| - n/k| \leq C\sqrt{n} \quad (68)$$

with probability at least 0.99. Writing $v \sim w$ if $vw \in E(G_n)$, for $1 \leq i, j \leq k$ we have

$$\begin{aligned} \mathbb{E}(e(P'_i, P'_j)) &= \sum_{(v,w) : v \sim w} \mathbb{E}(1_{i_v=i} 1_{i_w=j}) = \sum_{(v,w) : v \sim w} p_{v,i} p_{w,j} \\ &= n^2 p \int_{P_i \times P_j} \kappa_{G_n}(x, y) dx dy, \end{aligned}$$

so the expectation of $e(P'_i, P'_j)/(n^2 p)$ is exactly m_{ij}/k^2 . For edges $vw, v'w'$ of G_n , the random variables $1_{i_v=i} 1_{i_w=j}$ and $1_{i_{v'}=i} 1_{i_{w'}=j}$ are independent unless vw and $v'w'$ share a vertex, in which case their covariance is at most one. It follows that $\text{Var}(e(P'_i, P'_j))$ is bounded by $2 \text{hom}(P_2, G_n)$; the factor 2 arises

since we may put the common vertex of two incident edges into P_i or P_j . But $\text{hom}(P_2, G_n) \leq 2e(G_n)\Delta(G_n)$, which is $o(n^4p^2)$ by assumption. Hence, for any ε , the probability that we have

$$\left| \frac{e(P'_i, P'_j)}{n^2p} - \frac{m_{ij}}{k^2} \right| \leq \varepsilon \quad (69)$$

for every i and j with $1 \leq i, j \leq k$ is at least 0.99, provided n is large enough.

From the comments above, if n is large enough, there is a partition Π' for which both (68) and (69) hold. Starting from such a partition and moving at most $O(\sqrt{n}) = o(n)$ vertices of L (the set of low-degree vertices) between parts, we may find a balanced partition with almost the same density matrix. In other words, we may find an element of $\mathcal{M}_k(G_n)$ close to M , completing the proof. \square

If $np \rightarrow \infty$, then the condition of Lemma 6.1 that $\Delta(G_n) = o(n^2p)$ holds trivially, since $\Delta(G_n) \leq n = o(n^2p)$. When np is bounded, this condition is necessary. Taking G_n to be a star, for example, every partition of $V(G_n)$ has the property that there is one part meeting all edges. But the corresponding kernel has partitions which are very far from having this property, namely those in which, roughly speaking, the central vertex of the star has been split between parts.

6.2 The relationship between the cut and partition metrics

We now turn to the main result of this section, showing the equivalence of d_{cut} and d_{part} under mild assumptions. The key idea of the proof is that one can identify the density matrix corresponding to a weakly (ε, p) -regular partition from the set of density matrices.

Theorem 6.2. *Let $np \rightarrow \infty$, and let (G_n) be a sequence of graphs with $|G_n| = n$ satisfying the bounded density assumption 4.1. Let κ be a bounded kernel. Then $d_{\text{part}}(G_n, \kappa) \rightarrow 0$ if and only if $d_{\text{cut}}(G_n, \kappa) \rightarrow 0$.*

Proof. Suppose first that $d_{\text{cut}}(G_n, \kappa) \rightarrow 0$, i.e., that $d_{\text{cut}}(\kappa_{G_n}, \kappa) \rightarrow 0$. If κ_1 and κ_2 are any kernels with $d_{\text{cut}}(\kappa_1, \kappa_2) < d$, and $M \in \mathcal{M}_k(\kappa_1)$, then there is an $M' \in \mathcal{M}_k(\kappa_2)$ whose entries differ from those of M by at most k^2d : one simply takes the corresponding partition for κ_2 , after rearranging so that $\|\kappa_1 - \kappa_2\|_{\text{cut}} < d$. It follows that $d_H(\mathcal{M}_k(\kappa_1), \mathcal{M}_k(\kappa_2)) \leq k^2d_{\text{cut}}(\kappa_1, \kappa_2)$. Hence, $d_H(\mathcal{M}_k(\kappa_{G_n}), \mathcal{M}_k(\kappa)) \rightarrow 0$. Using Lemma 6.1, it follows that $d_{\text{part}}(G_n, \kappa) \rightarrow 0$.

Now suppose that $d_{\text{part}}(G_n, \kappa) \rightarrow 0$. By the *index* $\text{ind}(M)$ of a density matrix $M = (m_{ij}) \in \mathcal{M}_k(\kappa')$ we mean simply $k^{-2} \sum m_{ij}^2$. Let $f(k, \varepsilon) \geq k$ be a function to be specified later. A k -by- k density matrix $M \in \mathcal{M}_k(\kappa')$ is *locally ε -optimal* for a kernel κ' if

$$\sup_{\ell \leq f(k, \varepsilon)} \sup_{M' \in \mathcal{M}_\ell(\kappa')} \text{ind}(M') \leq \text{ind}(M) + \varepsilon,$$

i.e., if M has almost maximal index among density matrices with not too many parts; the definition of local optimality for $M \in \mathcal{M}_k(G_n)$ is similar.

Fix $\varepsilon > 0$. Since (G_n) has bounded density, whenever n is large enough as a function of k , any density matrix in $\mathcal{M}_k(G_n)$ has index at most some constant C . It follows that there is a $K = K(C, \varepsilon)$ such that, for n large enough, every G_n has some locally optimal density matrix $M_k(n)$ of size at most K . (This statement is a key part of the proof of Szemerédi's Lemma.)

Since $d_{\text{part}}(G_n, \kappa) \rightarrow 0$, if n is large enough, there is an $M'_k(n) \in \mathcal{M}_k(\kappa)$ with all entries within $\varepsilon/(10C)$ of those of $M_k(n)$. It follows that $\text{ind}(M'_k(n)) \geq \text{ind}(M_k(n)) - \varepsilon/2$. Similarly, for n large, every $M' \in \cup_{\ell \leq f(k, \varepsilon)} \mathcal{M}_\ell(\kappa)$ has all entries within $\varepsilon/(10C)$ of some $M \in \cup_{\ell \leq f(k, \varepsilon)} \mathcal{M}_\ell(G_n)$, which implies

$$\text{ind}(M') \leq \text{ind}(M) + \varepsilon/2 \leq \text{ind}(M_k(n)) + 3\varepsilon/2 \leq \text{ind}(M'_k(n)) + 2\varepsilon,$$

using the assumption that $M_k(n)$ is locally ε -optimal for G_n for the second inequality. Thus $M'_k(n)$ is locally 2ε -optimal for κ .

Recall that a partition Π of $[0, 1]$ is *weakly* (ε, p) -regular with respect to a kernel κ' if the corresponding averaged kernel κ'/Π satisfies $\|\kappa'/\Pi - \kappa'\|_{\text{cut}} \leq \varepsilon$. The proof of Lemma 4.3 (a sparse form of the Frieze-Kannan form of Szemerédi's Lemma) shows that if (G_n) has bounded density, then there is a function $f(k, \varepsilon)$ such that, if $n \geq n_0(k, \varepsilon)$ and $M \in \mathcal{M}_k(G_n)$ is locally ε -optimal, then the corresponding partition of κ_{G_n} is weakly (ε, p) -regular; the same applies to κ . It follows that for n large, identifying each density matrix with a corresponding kernel, we have $d_{\text{cut}}(\kappa_{G_n}, M_k(n))$, $d_{\text{cut}}(M_k(n), M'_k(n))$ and $d_{\text{cut}}(M'_k(n), \kappa)$ all of order $O(\varepsilon)$. Since ε was arbitrary, it follows that $d_{\text{cut}}(\kappa_{G_n}, \kappa) \rightarrow 0$, as required. \square

In the light of Corollary 4.7, Theorem 6.2 implies that a sequence (G_n) satisfying Assumption 4.1 is Cauchy with respect to d_{part} if and only if it is Cauchy with respect to d_{cut} .

The bounded density assumption in Theorem 6.2, which is trivially satisfied in the dense case $p = \Theta(1)$, is necessary in general. This can be seen by considering, for example, a graph G_n made up of n/m complete graphs of order m , with $m \sim pn = o(n)$ chosen so that G_n has $pn^2/2$ edges. By compactness, any sequence with $e(G_n) = O(pn^2)$ has a subsequence that is Cauchy with respect to d_{part} (here, in fact, the original sequence is Cauchy). However, it is easy to check that no subsequence of (G_n) is Cauchy with respect to d_{cut} .

The proof of Theorem 6.2 applies just as well to kernels as to graphs (and one can in any case approximate kernels by dense graphs), showing that $d_{\text{part}}(\kappa_n, \kappa) \rightarrow 0$ if and only if $d_{\text{cut}}(\kappa_n, \kappa) \rightarrow 0$. It follows that d_{part} induces a metric on \mathcal{K} , the set of kernels quotiented by equivalence, and that d_{part} and d_{cut} give rise to the same topology on \mathcal{K} . This was proved by Borgs, Chayes, Lovász, Sós and Vesztergombi [16] in their study of the dense case, as part of their Theorem 3.5.

7 Discussion and closing remarks

For dense graphs, with $\Theta(n^2)$ edges, the results of Borgs, Chayes, Lovász, Sós and Vesztegombi [15, 16] show that one single metric, say d_{cut} , effectively captures several natural notions of local and global similarity. Indeed, convergence in d_{cut} is equivalent to convergence in the partition metric d_{part} (a natural global notion) and to convergence in d_{sub} , i.e., convergence of all small subgraph counts, a natural local notion. These results apply to all sequences (G_n) of graphs, but if G_n has $o(n^2)$ edges then they become trivial: any such sequence is Cauchy with respect to any of the metrics, and indeed converges to the zero kernel. To make interesting statements about sparse graphs one should adapt the metrics so that, roughly speaking, given an ‘edge density function’ $p = p(n)$ satisfying $p \rightarrow 0$, one compares a graph G_n with $p\binom{n}{2}$ edges to the Erdős–Rényi random graph $G(n, p)$ and its inhomogeneous variants rather than to K_n . Our main aim in this paper has been to introduce such metrics, and to discuss the relationships between them. In this final section we turn to a slightly different question, that of the relationship between metrics and random graph models.

7.1 Models and metrics

In the dense case, there is a very natural correspondence between limit points of sequences converging in d_{cut} , and the inhomogeneous random graph model $G(n, \kappa)$. In general, given any metric, we can ask whether there is a corresponding random graph model: for each metric d on some class of (sparse) graphs satisfying certain restrictions, we can ask the following question.

Question 7.1. *Given a metric d , can we find a ‘natural’ family of random graph models with the following two properties: (i) for each model, the sequence of random graphs (G_n) generated by the model is Cauchy with respect to d with probability 1, and (ii) for any sequence (G_n) with $|G_n| = n$ that is Cauchy with respect to d , there is a model from the family such that, if we interleave (G_n) with a sequence of random graphs from the model, the resulting sequence is still Cauchy with probability 1.*

In the above question, we are implicitly assuming a coupling between the probability spaces on which the graphs (G_n) are defined. There is of course no need to do so: we can replace ‘Cauchy with probability 1’ with the less familiar ‘Cauchy in probability’, which is equivalent to convergence in probability in the completion; see Kallenberg [29, Lemma 4.6].

Although Question 7.1 is rather vague, for $d = d_{\text{cut}}$ the answer is ‘yes’ in the dense case, since (G_n) is Cauchy if and only if $d_{\text{cut}}(G_n, \kappa) \rightarrow 0$ for some kernel κ , while the dense inhomogeneous random graphs $G(n, \kappa)$ converge to κ in d_{cut} with probability 1. Thus our family consists of one model $G(n, \kappa)$ for each kernel κ (to be precise, for each equivalence class of kernels under the relation \sim defined in Subsection 2.4).

In the sparse case we do not have an entirely satisfactory answer for any of the metrics considered in this paper. Assuming that $np \rightarrow \infty$, there is an al-

most completely satisfactory answer for d_{cut} : if we impose the bounded density assumption 4.1, then Corollary 4.7 and Lemma 4.10 show that the sparse inhomogeneous models $G_p(n, \kappa)$ answer Question 7.1. For d_{sub} , defined with respect to certain restricted sets of subgraphs, the results in Section 5 (in particular, Theorem 5.18) show that once again $G_p(n, \kappa)$ answers this question for suitably restricted sequences.

The extremely sparse case, where $p = \Theta(1/n)$, turns out to be even more complicated; we shall discuss this in a forthcoming paper [11].

There is an even vaguer, but perhaps more important, ‘mirror image’ of Question 7.1. Suppose that we have a random graph model, and we would like to test whether it is appropriate for some network in the real world. Then we would like to have a suitable metric to compare a ‘typical’ graph from the model with the real-world network. It is too much to hope that one metric will be appropriate in all situations; in particular, taking the simple case in which our model is $G(n, p)$ for some $p = p(n) \rightarrow 0$, the unnormalized metrics d_{cut} , d_{sub} or d_{part} , that are very suitable for dense graphs, will declare any graph with $o(n^2)$ edges to be close to the model.

In general, a random graph model (or family of models) may suggest an appropriate metric, or at least properties such a metric should have. For example, the inhomogeneous models $G_p(n, \kappa)$ and the results here suggest the sparse version of d_{cut} . Suppose, however, that we are trying to model a network with rather few edges but high ‘clustering’, i.e., many triangles and other small subgraphs. One possible model is a denser version of the sparse random graphs with clustering introduced by Bollobás, Janson and Riordan [9]: given, for each fixed graph F , a ‘kernel’ $\kappa_F : [0, 1]^{|F|} \rightarrow [0, \infty)$ and a normalizing function $p_F(n)$, we choose vertex types x_1, \dots, x_n independently and uniformly at random and then, for each F , add each possible copy of F with vertex set v_1, \dots, v_k , $1 \leq v_1 < v_2 < \dots < v_k \leq n$, with probability $\kappa_F(x_{v_1}, \dots, x_{v_k})p_F(n)$.

In this model, a huge family of normalizations are possible: we can take each p_F to be any function of n bounded by 1. Of course, certain restrictions will be necessary for the model to make much sense; otherwise, for example, the copies of some F_1 added directly may be swamped by copies of F_1 arising as subgraphs of some F_2 , in which case there was no point adding any copies of F_1 directly. However, there is no doubt that many different normalizations will be interesting: for example, for any $0 < a \leq 4/3$, we can produce graphs with, say, $\Theta(n^{4/3})$ edges and $\Theta(n^a)$ triangles. Indeed, to do so we need only two kernels, one for edges (which we may take to generate a bipartite graph if needed), and one for triangles.

If, for some reason, we are considering graphs with, say, around $n^{4/3}$ edges and $n^{6/5}$ triangles, which is many more triangles than expected in $G(n, n^{-2/3})$, then the triangles are an important part of the structure, so in comparing two such graphs we should certainly compare the number of triangles, normalized by dividing by $n^{6/5}$. This suggests a family of metrics generalizing d_{sub} .

For each $F \in \mathcal{F}$ let $N_F = N_F(n)$ be a normalizing function satisfying $0 < N_F \leq \infty$. (We allow infinity to include the possibility of totally ignoring copies of some F . In fact, $N_F = n^{|F|+1}$ will do just as well.) Then we may define

a subgraph metric associated to $\mathbf{N} = (N_F)_{F \in \mathcal{F}}$ by modifying the definition of d_{sub} given in Section 3, using the normalized count $\text{emb}(F, G)/N_F(|G|)$ in place of $s_p(F, G)$. This metric will only make sense for suitably restricted families of graphs, but for such families, it will make much better sense than d_{sub} .

7.2 Closing Remarks

The main aim of this paper is to draw attention to the possibility that there is a rich theory of sparse (quasi-)random graphs waiting to be explored. The beginnings of such a theory can be found in the papers of Bollobás, Janson and Riordan [8, 9] in the very sparse case, and of Borgs, Chayes, Lovász, Sós, Szegedy and Vesztergombi [13, 14, 34] in the dense case; it would be desirable to build a theory encompassing these two extreme threads. As we have just shown, this task is unlikely to be easy: there are numerous unexpected difficulties and pitfalls, and much work has to be done even to arrive at concrete problems whose solutions would represent genuine progress in this endeavour. In this paper we have attempted to do some of this groundwork, and have identified some intriguing problems.

Our main focus has been the introduction of normalized versions of the metrics d_{cut} , d_{sub} and d_{part} , adapted to the study of graphs with $\Theta(pn^2)$ edges, where $p = p(n) \rightarrow 0$. We have shown in Section 6 that (under a mild assumption) d_{cut} and d_{part} have the same Cauchy sequences, and in Section 4 that (again under a mild assumption) these metrics have the property that any sequence (G_n) contains a subsequence converging to a kernel.

Turning to d_{sub} , things become more difficult. We have conjectured that if our p -normalized subgraph counts are suitably bounded and $p = p(n)$ is not too small then an appropriate Cauchy sequence does converge to a kernel (see Conjectures 3.3 and 3.4). Tantalizingly, we cannot even prove this convergence in just about the simplest case, when we know that the limit *has to be* a constant kernel (Conjecture 3.9).

Section 5 is devoted to the relationship between d_{cut} and d_{sub} . A sound understanding of the relationship between these two metrics, the cut and count metrics, would bring us much closer to a proper theory of sparse inhomogeneous quasi-random graphs. We have conjectured that under some natural and not too restrictive conditions, these two metrics are equivalent in the sense that if (G_n) is a sequence of graphs that are not too ‘lumpy’ then (G_n) converges to a kernel κ in the p -cut metric if and only if it converges to κ in the p -count metric (see Conjecture 5.6). As one of our main results, we have proved that p -cut convergence does imply p -count convergence for a restricted set of subgraph counts, under a mild assumption on the distribution of paths of certain lengths (see Theorems 5.15 and 5.17).

The case of graphs of bounded average degree turns out to be even more difficult, and will be discussed in a companion paper [11].

Acknowledgement. The authors would like to thank an anonymous referee for many detailed suggestions improving the presentation of the paper.

References

- [1] N. Alon, Explicit Ramsey graphs and orthonormal labelings, *Elec. J. Combin.* **1** (1994), R12 (electronic), 8 pp.
- [2] G.R. Blakley and P. Roy, A Hölder type inequality for symmetric matrices with nonnegative entries, *Proc. Amer. Math. Soc.* **16** (1965), 1244–1245.
- [3] B. Bollobás, Threshold functions for small subgraphs, *Math. Proc. Cam. Phil. Soc* **90** (1981), 197–206.
- [4] B. Bollobás, *Modern Graph Theory*, Graduate Texts in Mathematics, vol. 184, Springer, New York (1998), xiv + 394 pp.
- [5] B. Bollobás, *Linear Analysis*, 2nd ed., Cambridge University Press (1999), xii + 240 pp.
- [6] B. Bollobás, *Random Graphs*, 2nd ed., Cambridge University Press (2001), xviii + 498 pp.
- [7] B. Bollobás, C. Borgs, J.T. Chayes and O. Riordan, Percolation on dense graph sequences, preprint (2007) (revised 2008). [arXiv:0701346](#).
- [8] B. Bollobás, S. Janson and O. Riordan, The phase transition in inhomogeneous random graphs, *Random Structures and Algorithms*, **31** (2007), 3–122.
- [9] B. Bollobás, S. Janson and O. Riordan, Sparse random graphs with clustering, preprint (2008). [arXiv:0807.2040](#).
- [10] B. Bollobás, S. Janson and O. Riordan, The cut metric, random graphs, and branching process, preprint (2009). [arXiv:0901.2091](#).
- [11] B. Bollobás and O. Riordan, Sparse graphs: metrics and random models, preprint (2008). [arXiv:0812.2656](#).
- [12] C. Borgs, J.T. Chayes and L. Lovász, Moments of two-variable functions and the uniqueness of graph limits, preprint (2008). [arXiv:0803.1244](#).
- [13] C. Borgs, J.T. Chayes, L. Lovász, V.T. Sós and K. Vesztergombi, Counting graph homomorphisms, in *Topics in Discrete Mathematics* (eds. M. Klazar, J. Kratochvil, M. Loeb, J. Matousek, R. Thomas, P. Valtr), Springer (2006), pp 315–371.
- [14] C. Borgs, J.T. Chayes, L. Lovász, V.T. Sós, B. Szegedy and K. Vesztergombi, Graph limits and parameter testing, *Proc. 38th ACM Symp. Theory of Computing* (2006), 261–270.
- [15] C. Borgs, J.T. Chayes, L. Lovász, V.T. Sós and K. Vesztergombi, Convergent sequences of dense graphs I: Subgraph frequencies, metric properties and testing, preprint (2006) (revised 2007). [arXiv:0702004](#).

- [16] C. Borgs, J.T. Chayes, L. Lovász, V.T. Sós and K. Vesztergombi, Convergent sequences of dense graphs II: Multiway cuts and statistical physics, preprint (2007). <http://www.cs.elte.hu/~lovasz/ConvRight.pdf>.
- [17] F. Chung and R. Graham, Sparse quasi-random graphs, *Combinatorica* **22** (2002), 217–244.
- [18] F.R.K. Chung, R.L. Graham and R.M. Wilson, Quasi-random graphs, *Combinatorica* **9** (1989), 345–362.
- [19] P. Diaconis and S. Janson, Graph limits and exchangeable random graphs, *Rendiconti di Matematica* **28** (2008), 33–61.
- [20] J. Dugundji, *Topology*, Allyn and Bacon, Inc., Boston, Mass. (1966), xvi+447 pp.
- [21] P. Erdős and A. Rényi, On random graphs. I, *Publ. Math. Debrecen* **6** (1959), 290–297.
- [22] P. Erdős and A. Rényi, On a problem in the theory of graphs, *Magyar Tud. Akad. Mat. Kutató Int. Közl.* **7** (1962), 623–641.
- [23] A. Frieze and R. Kannan, Quick approximation to matrices and applications, *Combinatorica* **19** (1999), 175–220.
- [24] S. Gerke and A. Steger, The sparse regularity lemma and its applications, in *Surveys in combinatorics 2005*, London Math. Soc. Lecture Notes **327**, Cambridge University Press (2005), pp. 227–258.
- [25] E.N. Gilbert, Random graphs, *Ann. Math. Statist.* **30** (1959), 1141–1144.
- [26] S. Janson, Poisson approximation for large deviations, *Random Structures and Algorithms* **1** (1990), 221–230.
- [27] S. Janson, Standard representation of multivariate functions on a general probability space, preprint (2008).arXiv:0801.0196.
- [28] S. Janson, K. Oleszkiewicz and A. Ruciński, Upper tails for subgraph counts in random graphs, *Israel J. Math.* **142** (2004), 61–92.
- [29] O. Kallenberg, *Foundations of Modern Probability*, 2nd ed., Springer, New York, 2002.
- [30] J.H. Kim, The Ramsey number $R(3, t)$ has order of magnitude $t^2/\log t$, *Random Structures Algorithms* **7** (1995), 173–207.
- [31] Y. Kohayakawa, Szemerédi’s regularity lemma for sparse graphs, in *Foundations of computational mathematics (Rio de Janeiro, 1997)*, Springer (1997), pp. 216–230.

- [32] Y. Kohayakawa and V. Rödl, Szemerédi's regularity lemma and quasi-randomness, in *Recent advances in algorithms and combinatorics*, CMS Books Math. **11**, Springer (2003), pp. 289–351
- [33] L. Lovász and V.T. Sós, Generalized quasirandom graphs, *J. Combin. Theory B* **98** (2008), 146–163.
- [34] L. Lovász and B. Szegedy, Limits of dense graph sequences, *J. Combin. Theory B* **96** (2006), 933–957.
- [35] A. Ruciński, When are small subgraphs of a random graph normally distributed? *Probab. Th. and Related Fields* **78** (1988), 1–10.
- [36] E. Szemerédi, Regular partitions of graphs, in *Problèmes combinatoires et théorie des graphes* (Colloq. Internat. CNRS, Univ. Orsay, Orsay, 1976), *Colloq. Internat. CNRS* **260**, CNRS, Paris (1978), pp. 399–401.
- [37] A. Thomason, Pseudo-random graphs, in *Proceedings of Random Graphs* (M. Karonski, ed.), Poznań, 1985, *Annals of Discrete Mathematics*, **33** (1987) 307–331.
- [38] A. Thomason, Random graphs, strongly regular graphs and pseudorandom graphs, in *Surveys in Combinatorics 1987*, London Math. Soc. Lecture Note Ser. **123**, Cambridge Univ. Press, Cambridge (1987), pp 173–195.